

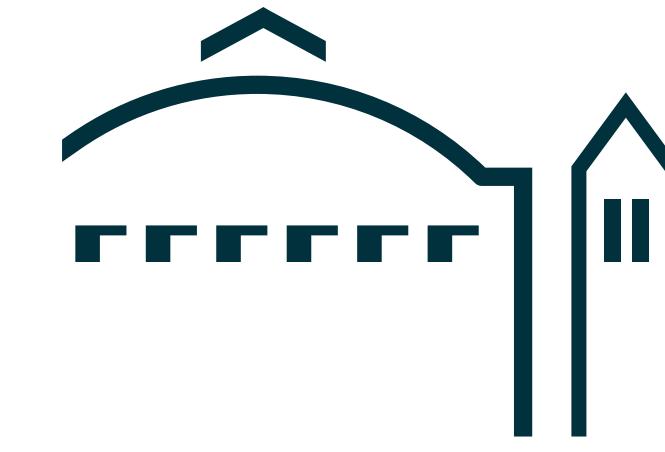
Expect the Unexpected

Lessons from a Decade of Passive Network Measurements

Johanna Amann
johanna@icir.org

<http://www.icir.org/johanna>

Who am I?



Expect the Unexpected

Interesting network phenomena and problems
you cannot typically write a paper about

Johanna Amann
johanna@icir.org

<http://www.icir.org/johanna>

Part 1

Passive Network Measurements



Enigma K

Enigma K est une variante de l'Enigma A utilisée par les forces armées britanniques et canadiennes pendant la Seconde Guerre mondiale. Elle est équipée d'un système de codage supplémentaire qui ajoute une autre couche de sécurité à l'algorithme standard. La machine est également munie d'un système de synchronisation des rotors supplémentaires.

La machine Enigma K a été modifiée spécialement pour les besoins militaires. Un éclairage supplémentaire a été ajouté. Une alimentation séparée est fournie pour les rotors supplémentaires.

Le chiffrement utilise 4 rotors et chaque rotor a 25 positions correspondant à l'alphabet. La clé est définie par 4 lettres, déterminant la position initiale des 4 rotors. Il y a donc $4! = 4 \times 3 \times 2 \times 1 = 24$ clés possibles pour ces 4 rotors. Les 3 rotors à droite peuvent être échangés entre eux (c'est un « miroir ») et ne peut pas être échangé avec les deux premiers rotors.

Un rotor est défini par une permutation sur l'alphabet. Chaque rotor contient une autre permutation. Les rotors sont placés dans l'appareil de manière à ce que leur position change à chaque tour de fonctionnement. Les opérations suivent maintenant 6 positions possibles : 6 + 7 + 2 + 1.

Enigma K

Enigma K est une variante de l'Enigma A utilisée par les forces armées britanniques et canadiennes pendant la Seconde Guerre mondiale. Elle est équipée d'un système de codage supplémentaire qui ajoute une autre couche de sécurité à l'algorithme standard. La machine est également munie d'un système de synchronisation des rotors supplémentaires.

Enigma K



2011

Apple Introduces iCloud

Free Cloud Services Beyond Anything Offered to Date

SAN FRANCISCO—June 6, 2011—Apple® today introduced iCloud®, a breakthrough set of free new cloud services that work seamlessly with applications on your iPhone®, iPad®, iPod touch®, Mac® or PC to automatically and wirelessly store your content in iCloud and automatically and wirelessly push it to all your devices. When anything changes on one of your devices, all of your devices are wirelessly updated almost instantly.

Apple Introduces iCloud

Free Cloud Service

SAN FRANCISCO—Just a breakthrough set of free services with applications on your iPhone automatically and wirelessly sync changes on one of your devices, updated almost instantl



Official Blog

Insights from Googlers into our products, technology, and the Google culture

Introducing the Google+ project: Real-life sharing, rethought for the web

June 28, 2011

Update: For our international readers, this post is also available in [French](#), [German](#), [Italian](#), [Japanese](#), [Portuguese](#), [Russian](#) and [Spanish](#). - Ed.

Among the most basic of human needs is the need to connect with others. With a smile, a laugh, a whisper or a cheer, we connect with others every single day.

Apple Introduces iCloud

Free Cloud Service

SAN FRANCISCO—Just a breakthrough set of free services with applications on your iPhone automatically and wirelessly, automatically and wirelessly, changes on one of your devices updated almost instantl

Google

Official Blog

Microsoft acquires Skype



Introducing the
for the web

June 28, 2011

Update: For our international readers, we've added support for French, German, Italian, Japanese, Portuguese, Spanish, and more.

Among the most basic things you can do in a video call is share a smile, a laugh, a whisper, or a kiss. Now you can do all of those things and more with your friends and family via video calls on the web.



Microsoft purchased Skype for \$8.5 billion, deepening the company's longstanding focus on real-time video and voice communications, and providing new market opportunities serving Skype's 160-plus million active users

Apple Introduces iCloud

Free Cloud Service

SAN FRANCISCO—Just a breakthrough set of features with applications on your iPhone automatically and wirelessly changes on one of your devices updated almost instantl

Google

Official Blog

Microsoft acquires Skype

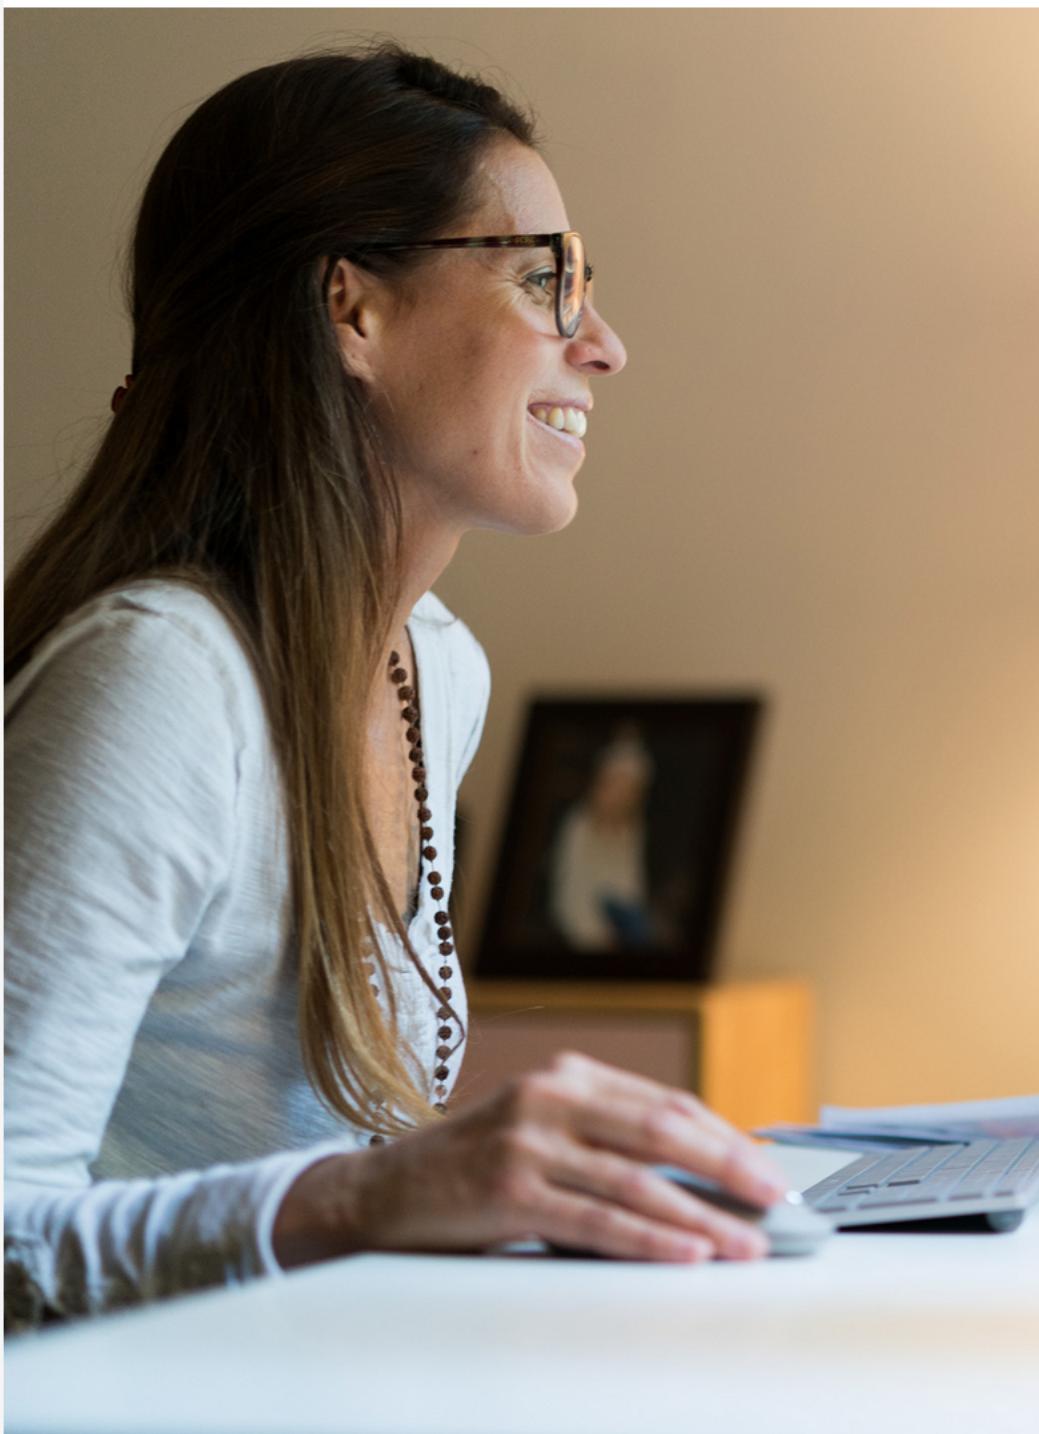


Introducing the
for the web

June 28, 2011

Update: For our international readers: We now have versions in Italian, Japanese, Portuguese, and Spanish.

Among the most basic things you can do is type, smile, a laugh, a whisper, or a kiss.



Microsoft purchased Skype for \$8.5 billion, deepening the company's focus on voice communications, and providing new market opport

Spotify

This article is more than 10 years old

Spotify announces launch in the US

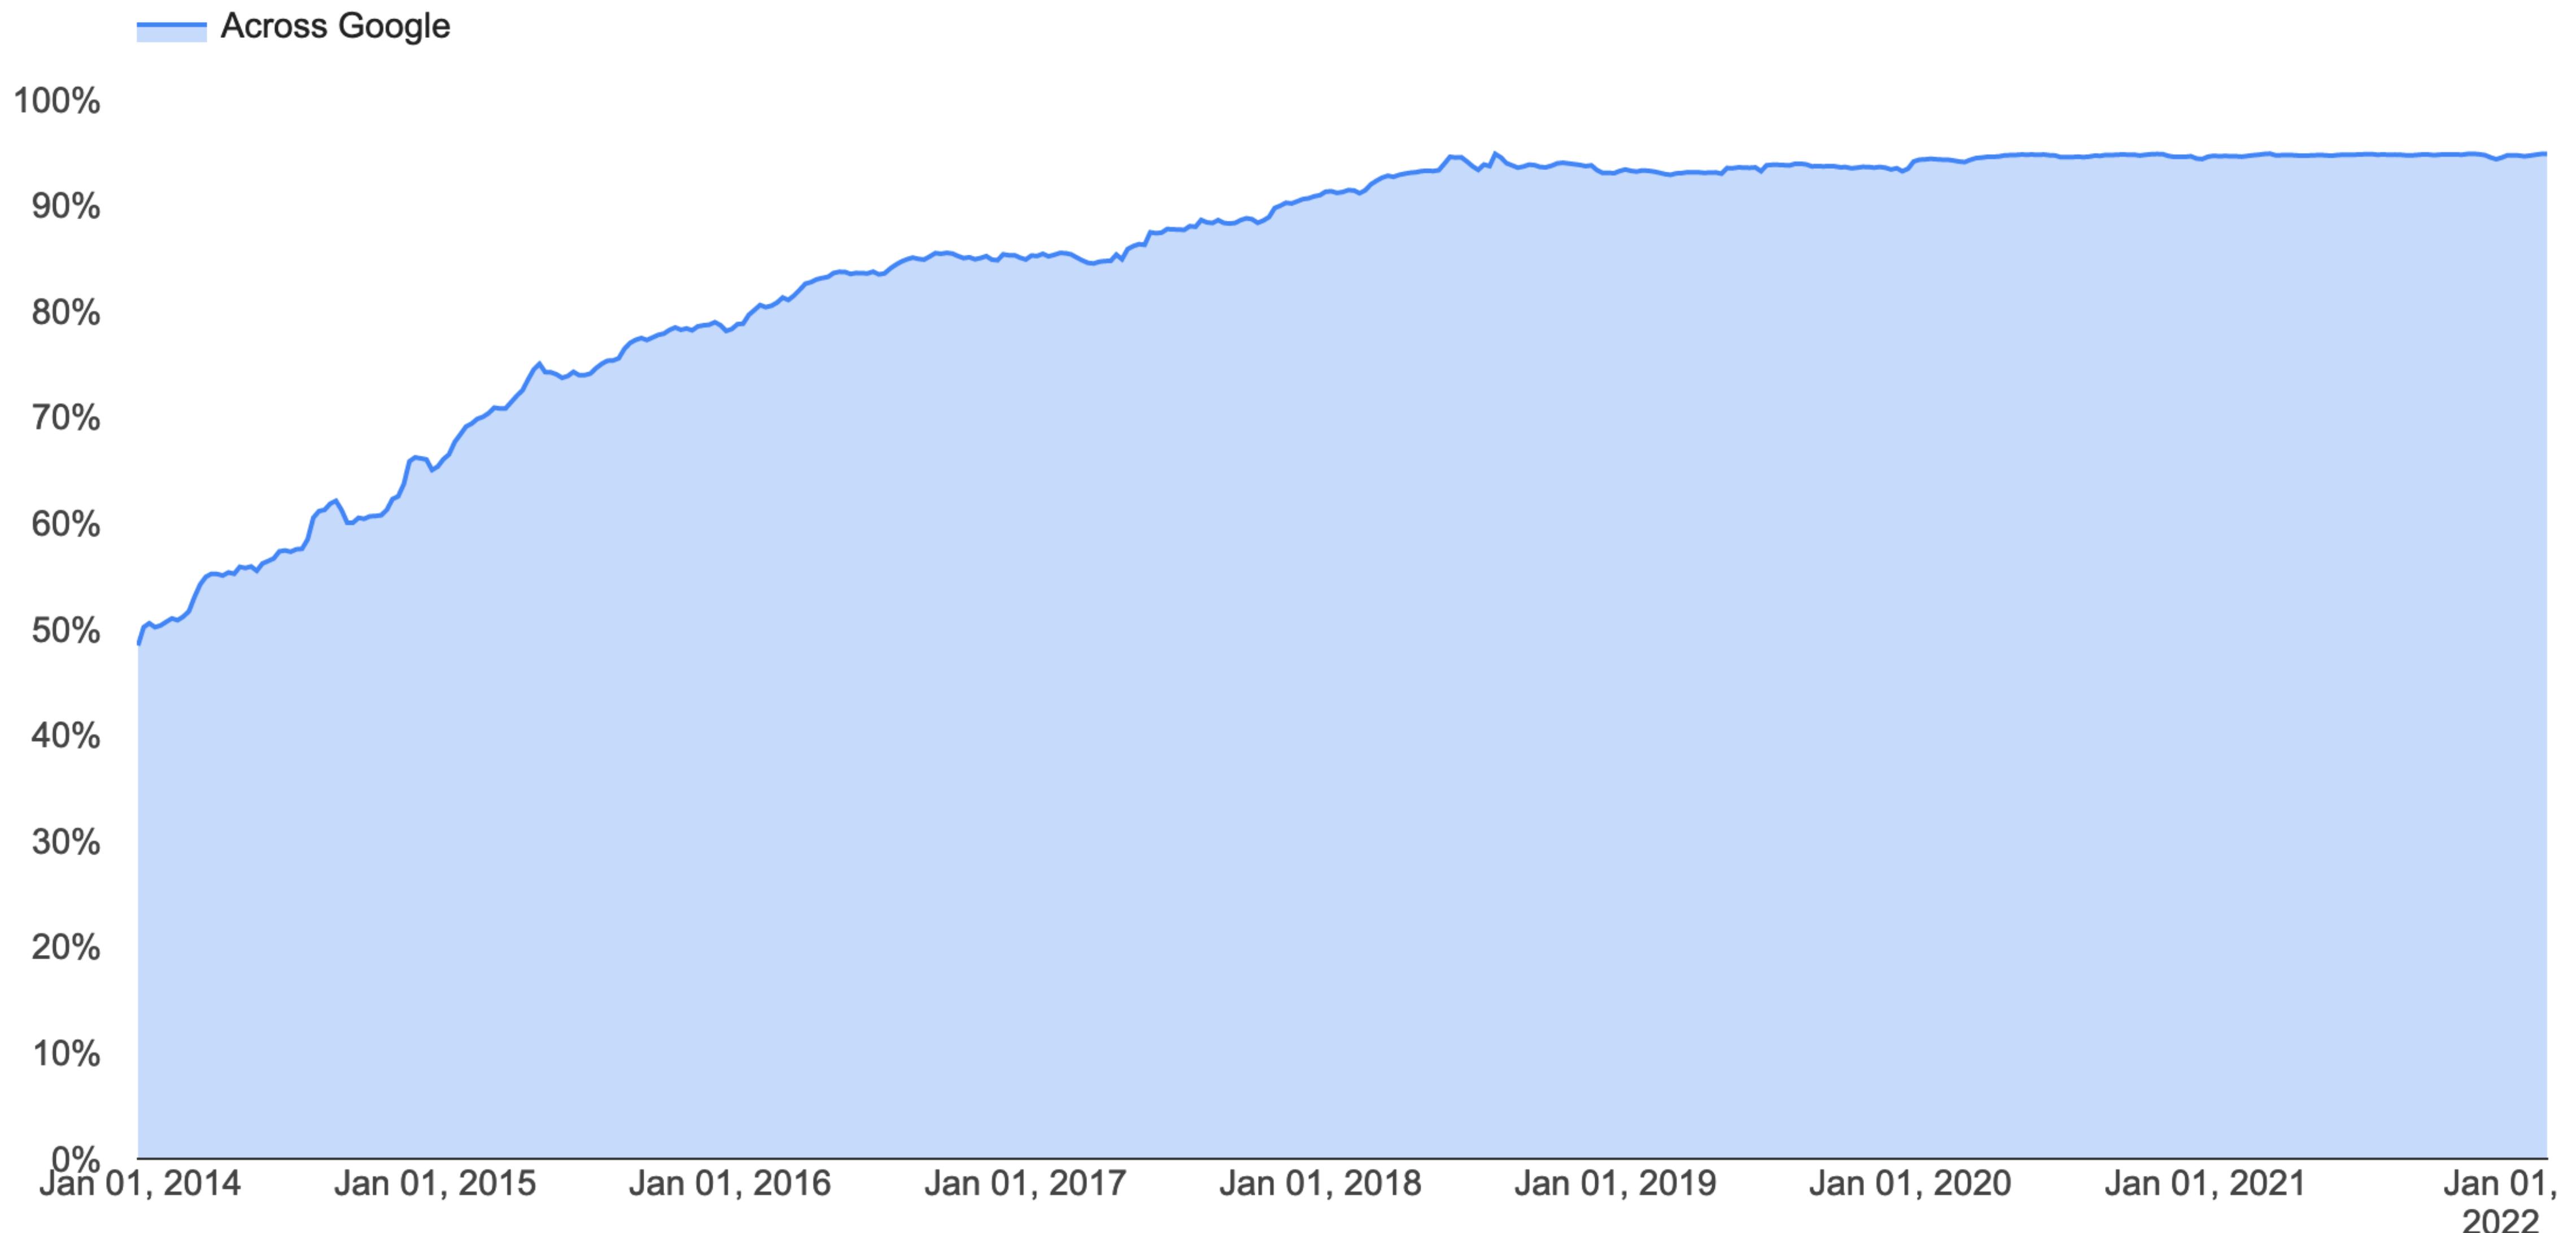
Swedish music streaming service expands beyond Europe for first time after two years of negotiations with record companies



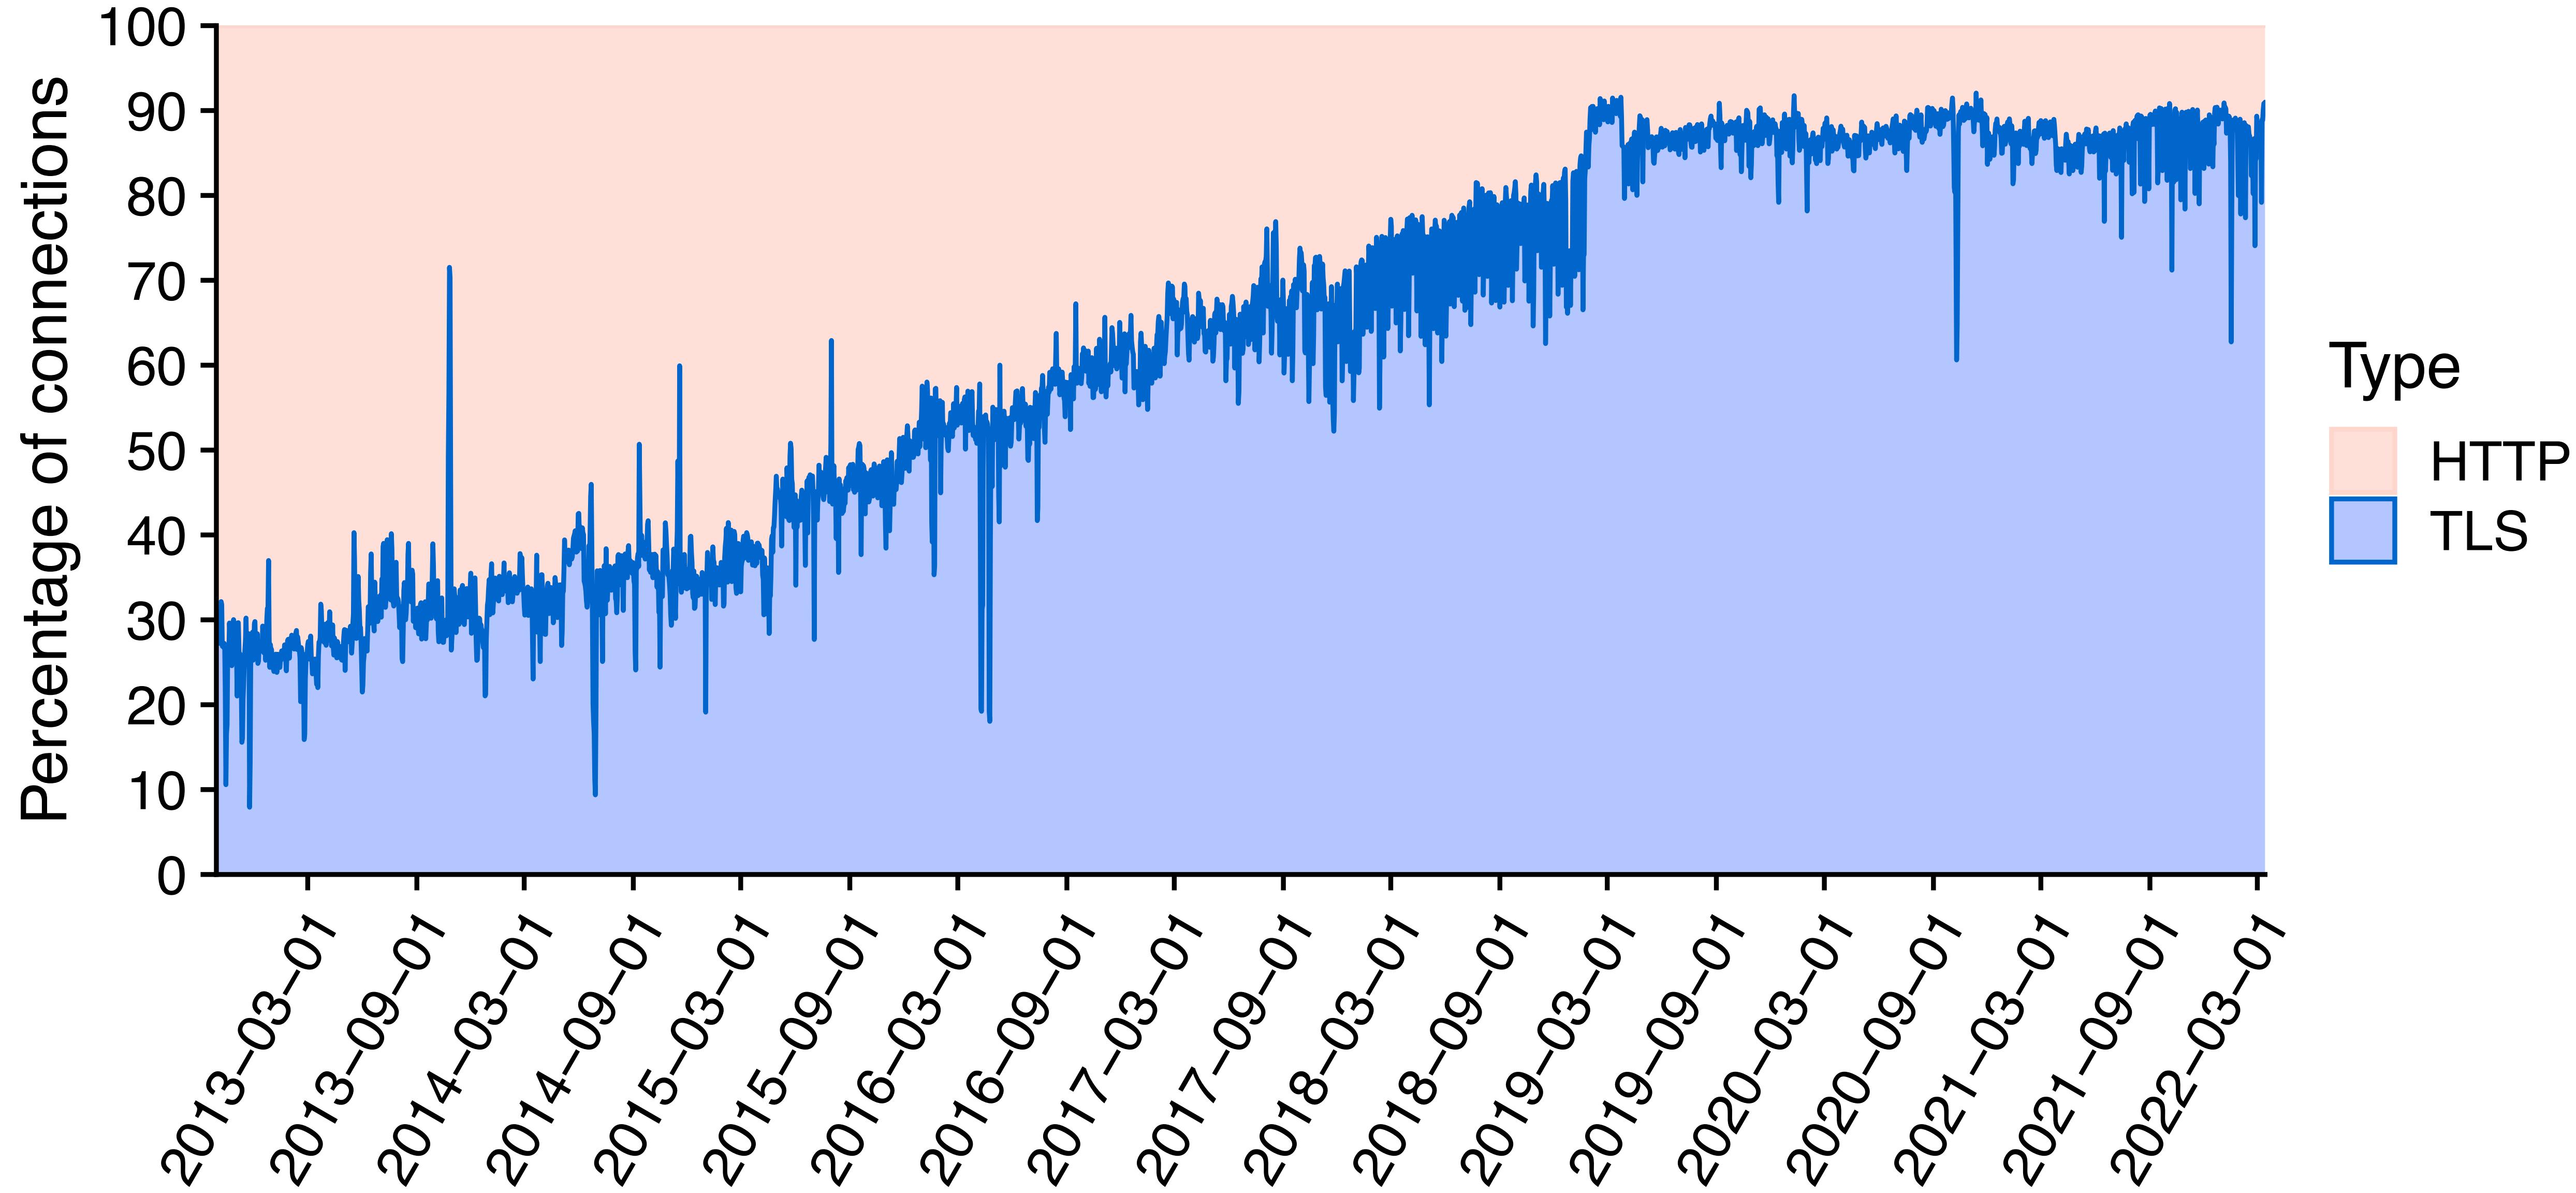
© Daniel Ek, co-founder and CEO of Spotify: the company has been valued at around \$1bn. Photograph: Rasmus Andersson/Spotify

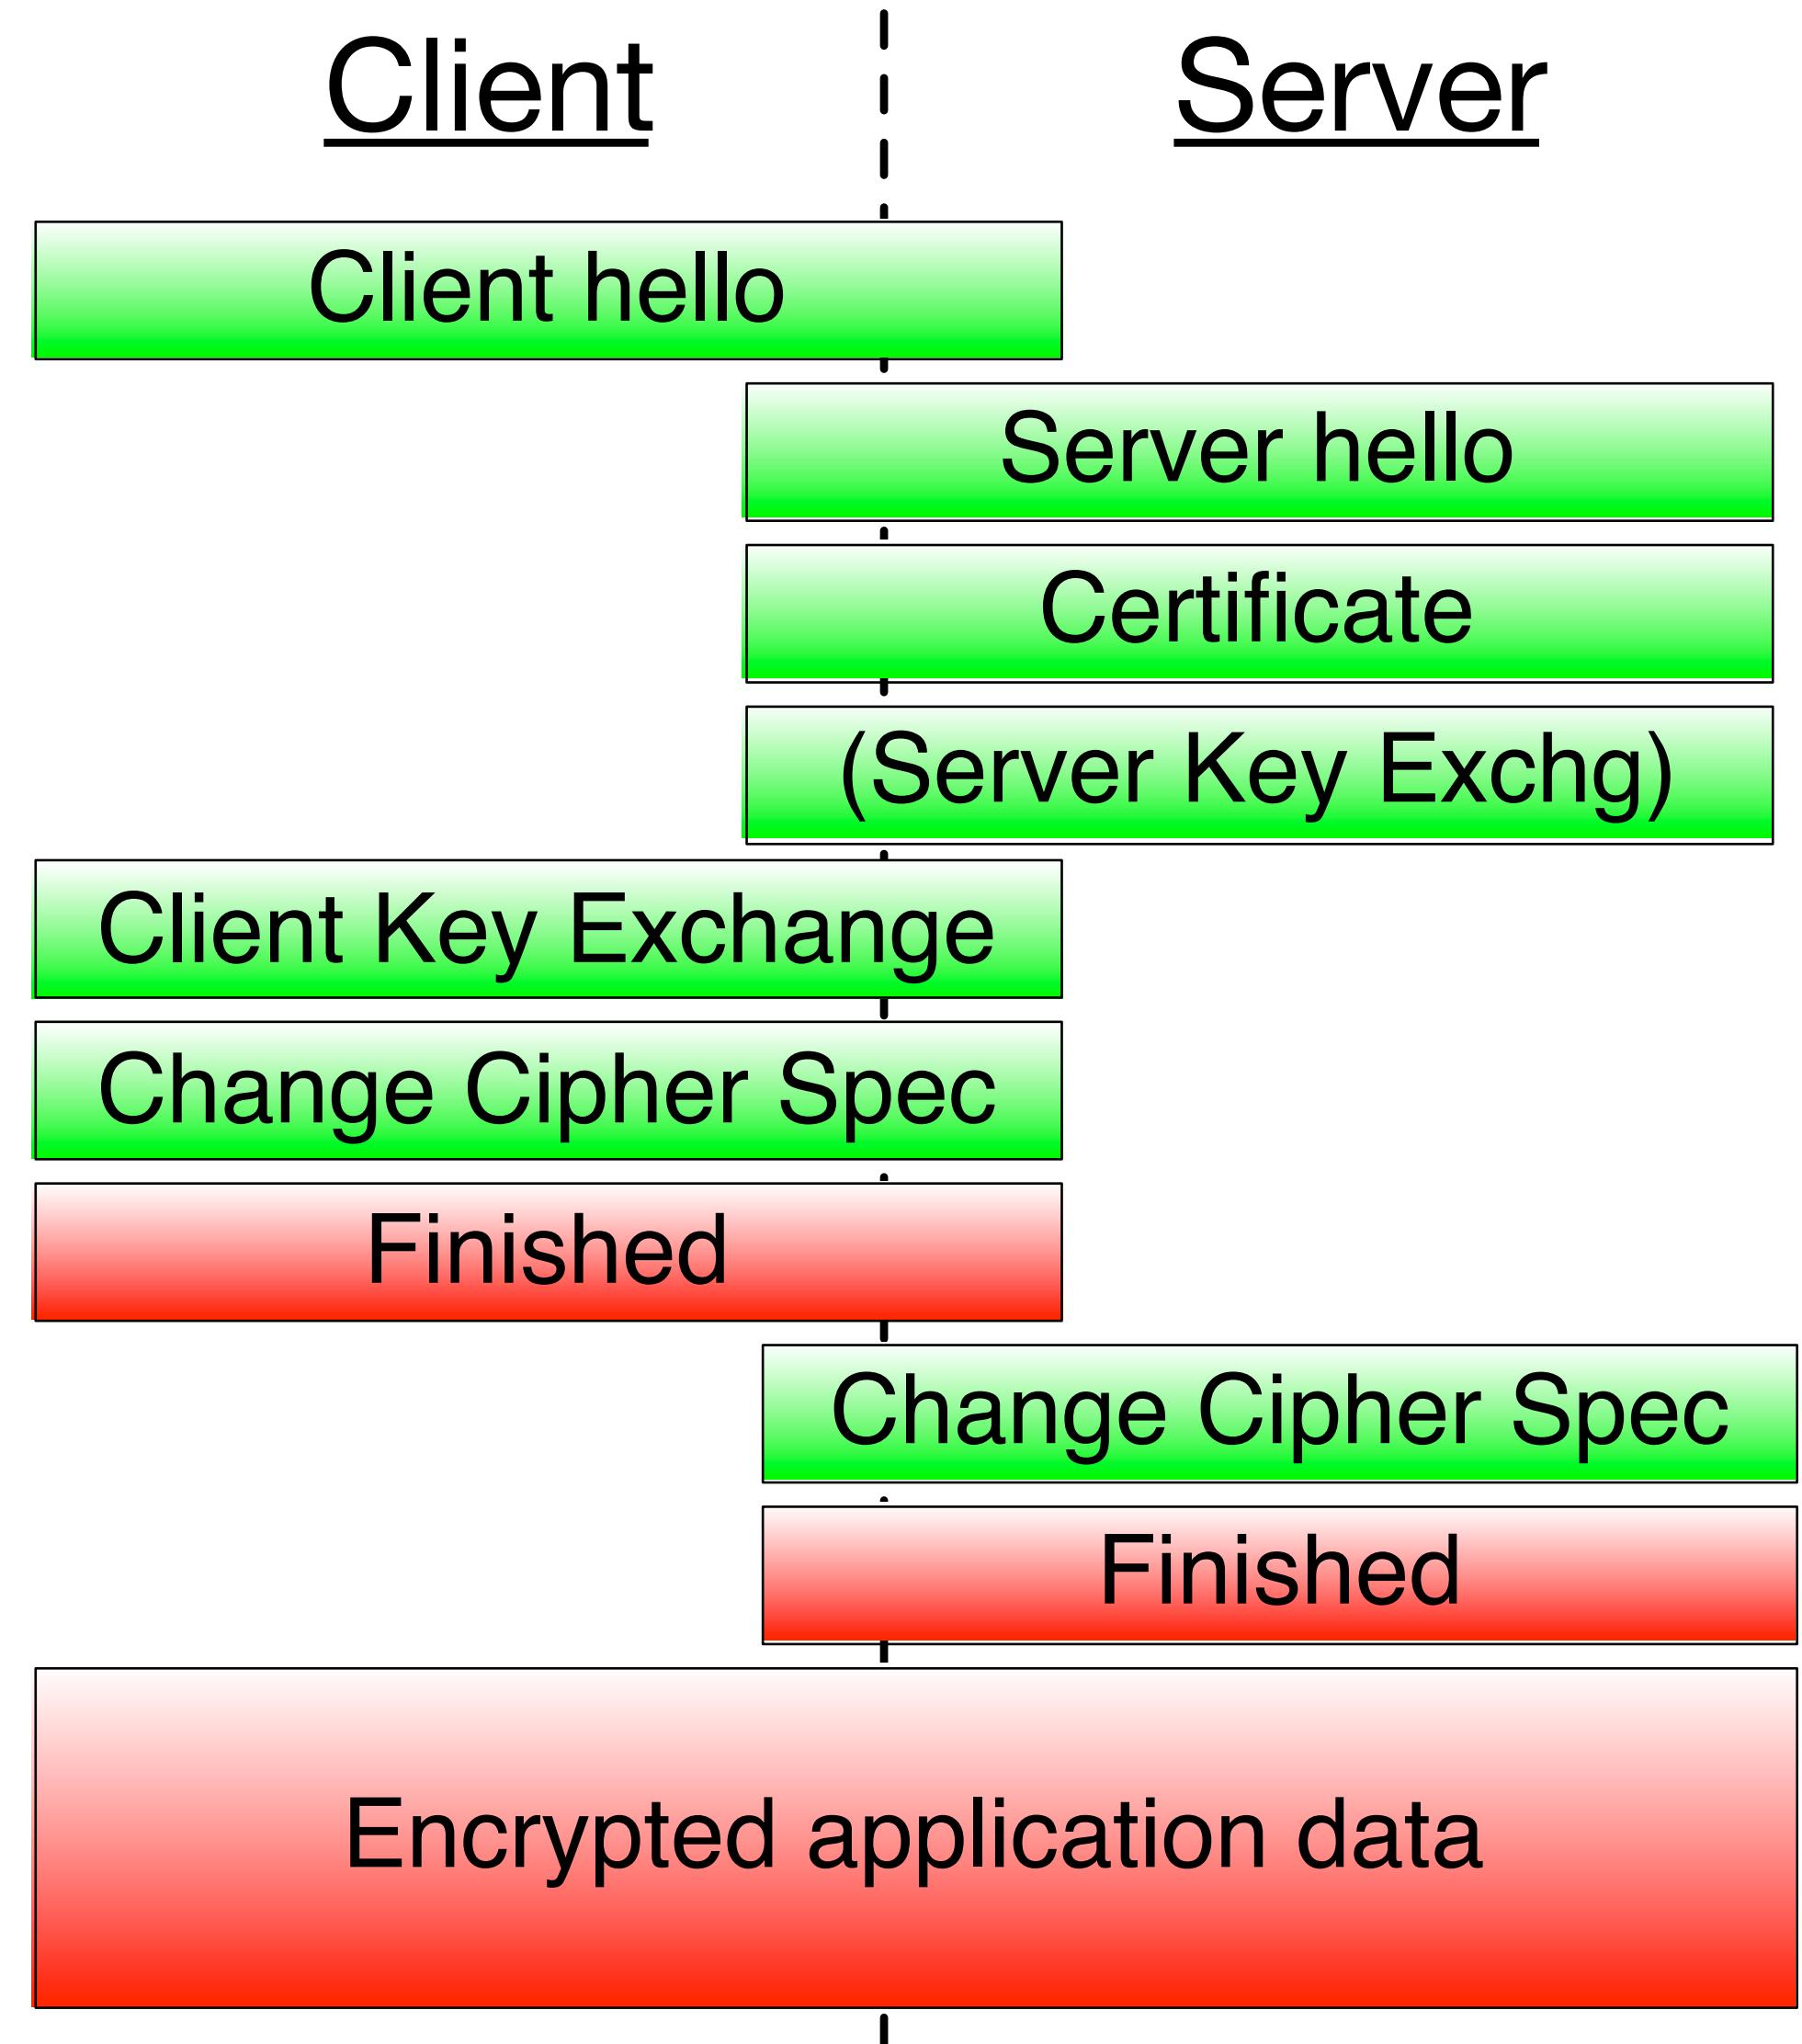
It has taken two years of intense wrangling with record companies and triggered wave after wave of speculation about future business models for the industry - now Spotify, the Swedish streaming music service, has finally announced that it will launch in the world's biggest music market, the US, on Thursday.

Encrypted Traffic across Google



Encrypted traffic at Berkeley Lab





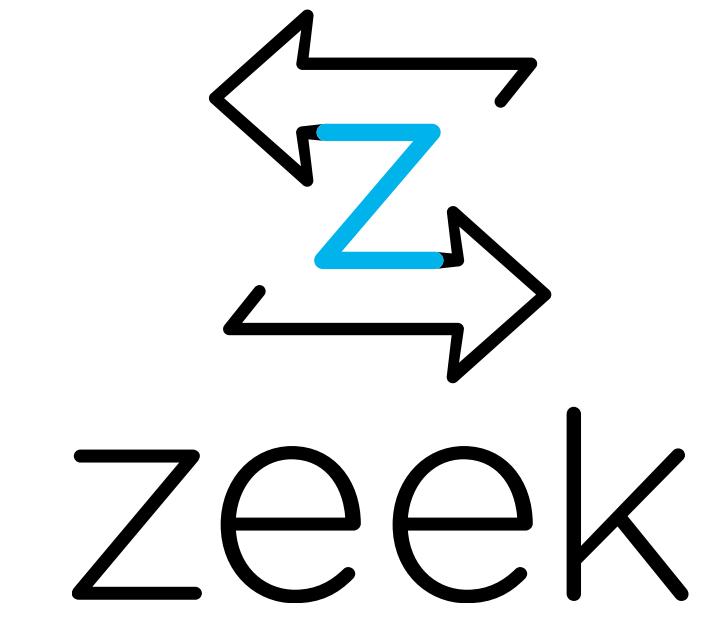
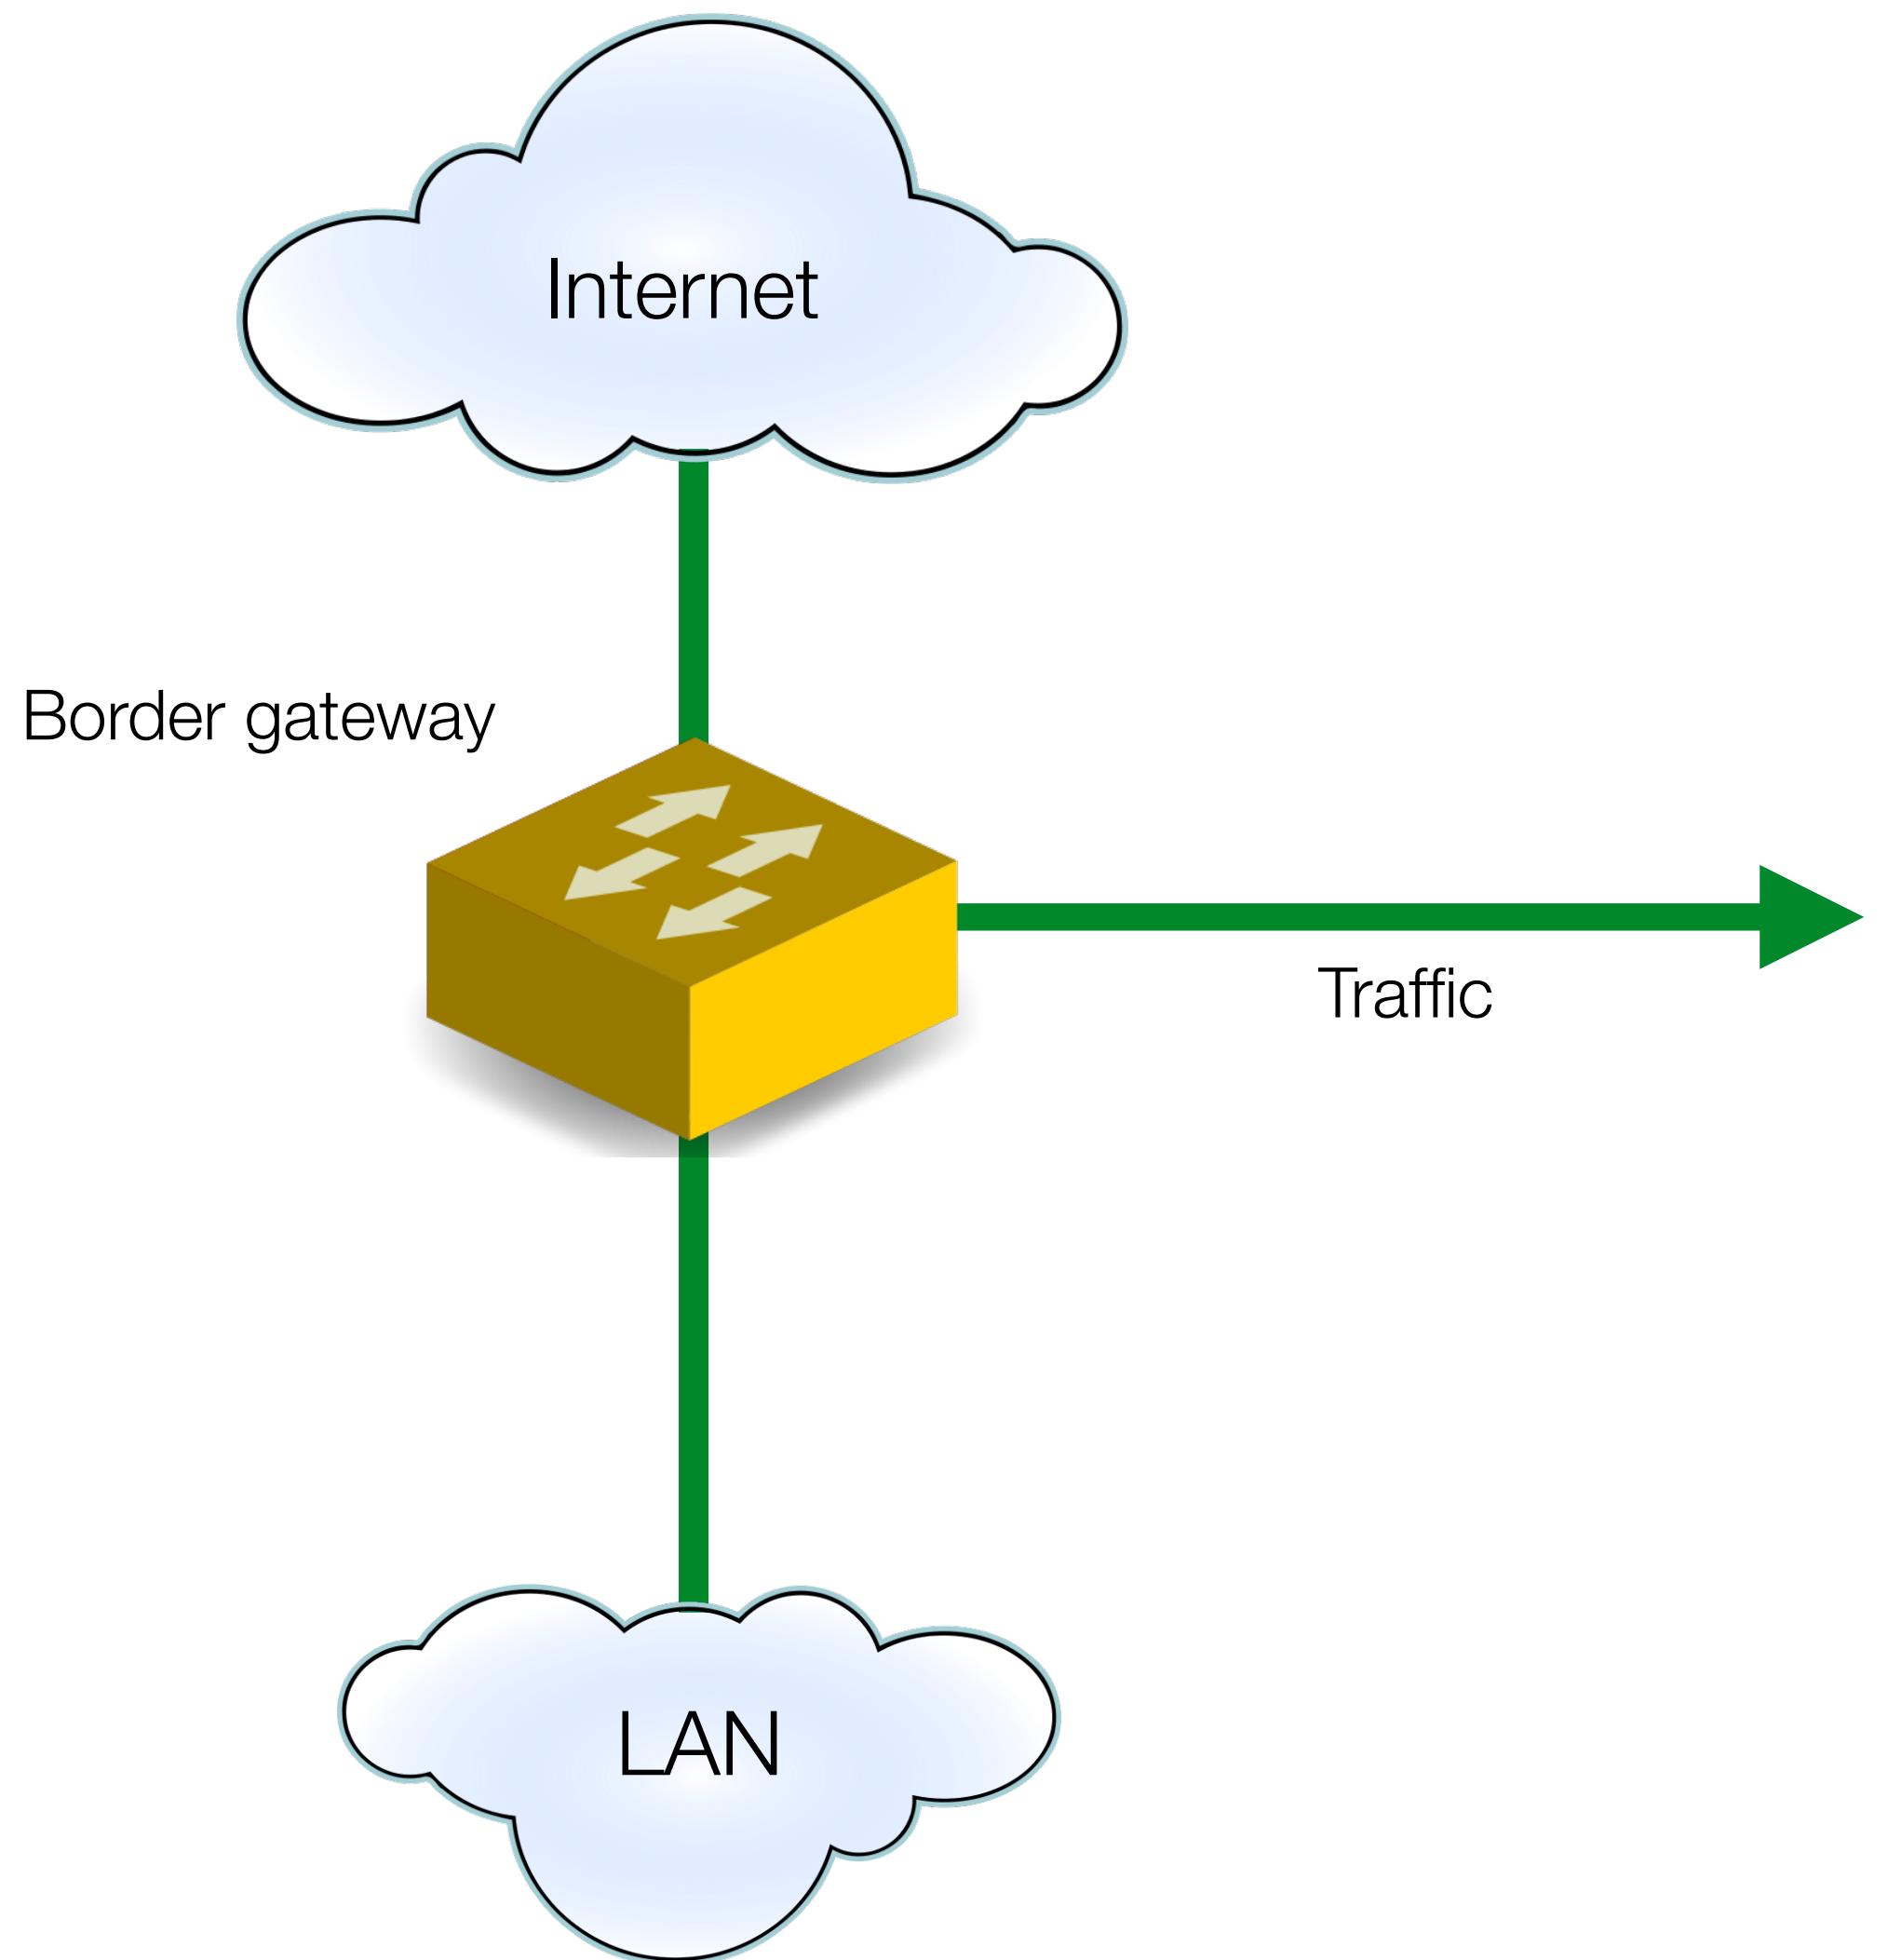


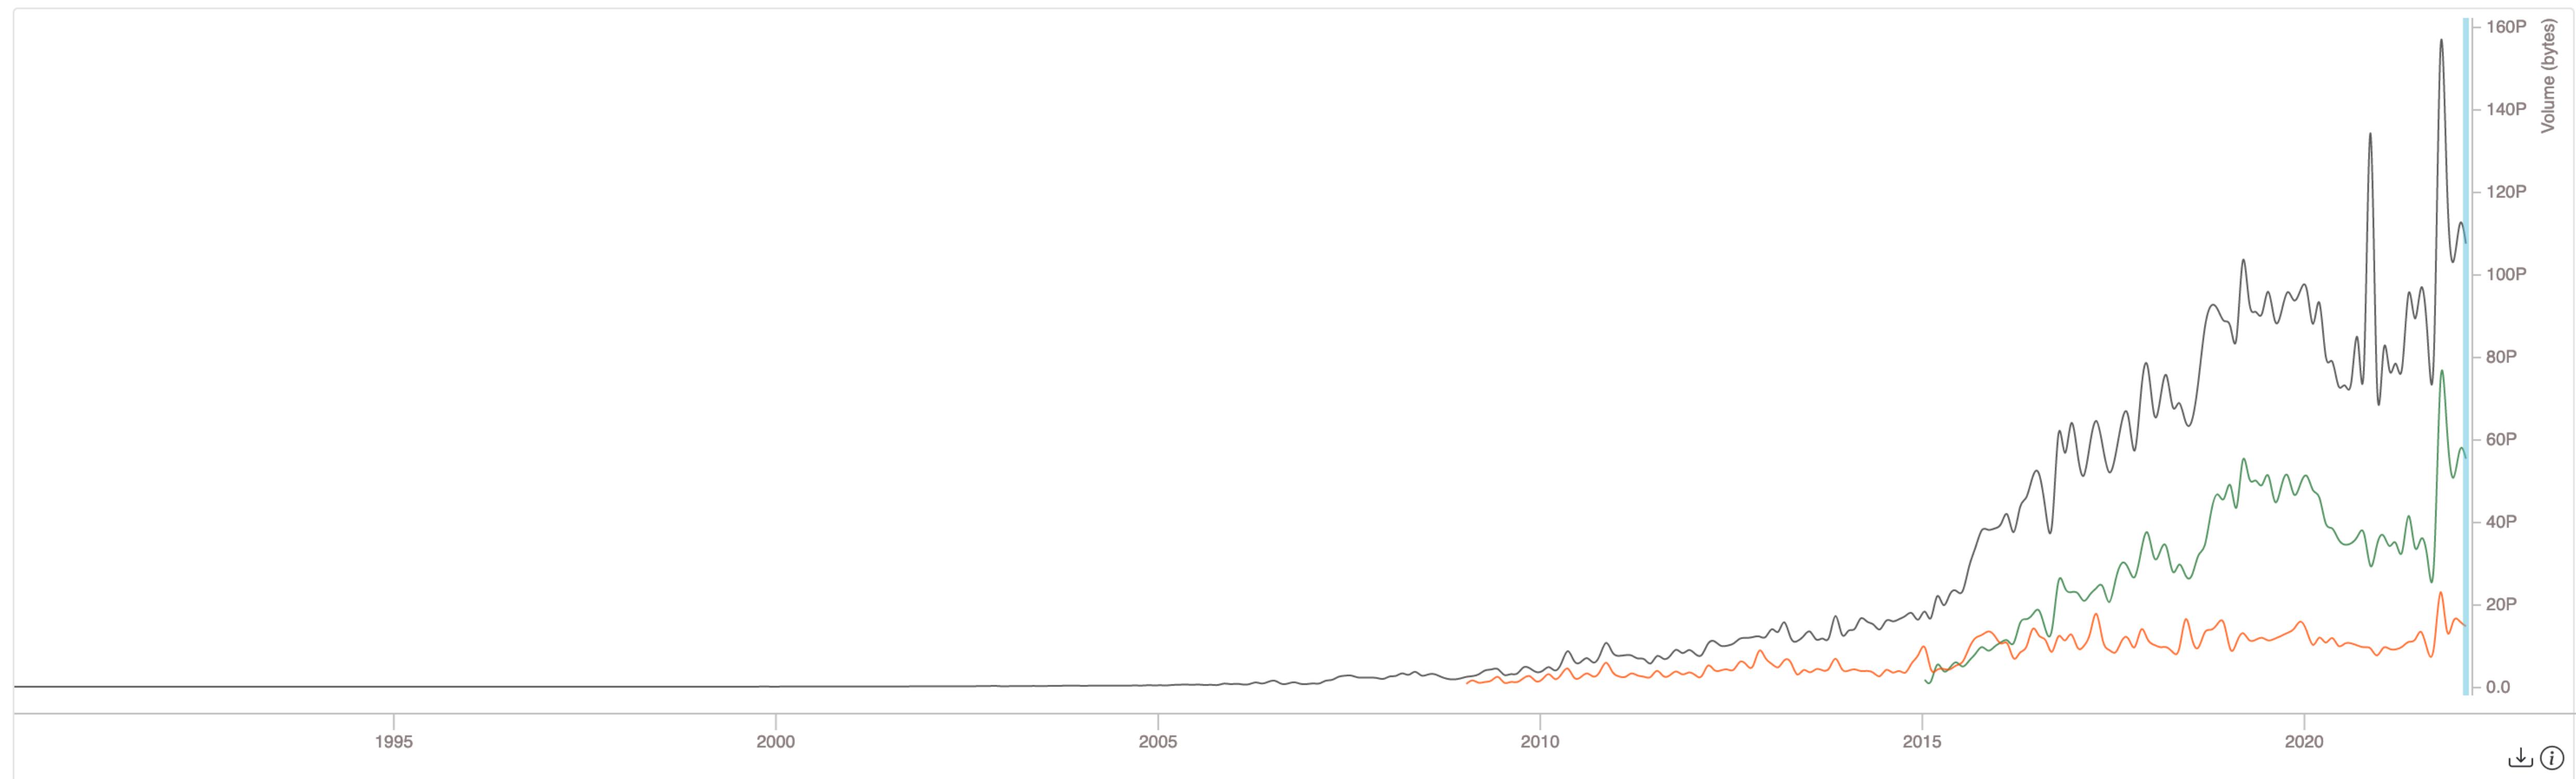
87 88 89 90 91 92 93 94 95 96 97 98
100 101 102 103 104 105 106 107 108 109 110 111
112 113 114 115 116 117 118
122 123 124 125 126 127 128 129
133 134 135 136 137 138 139 140 141
142 143 144 145 146 147 148 149 150
151 152 153 154 155 156 157 158 159
160 161 162 163 164 165 166 167 168
169 170 171 172 173 174 175 176 177
178 179 180 181 182 183 184 185 186
187 188 189 190 191 192 193 194 195
196 197 198 199 200

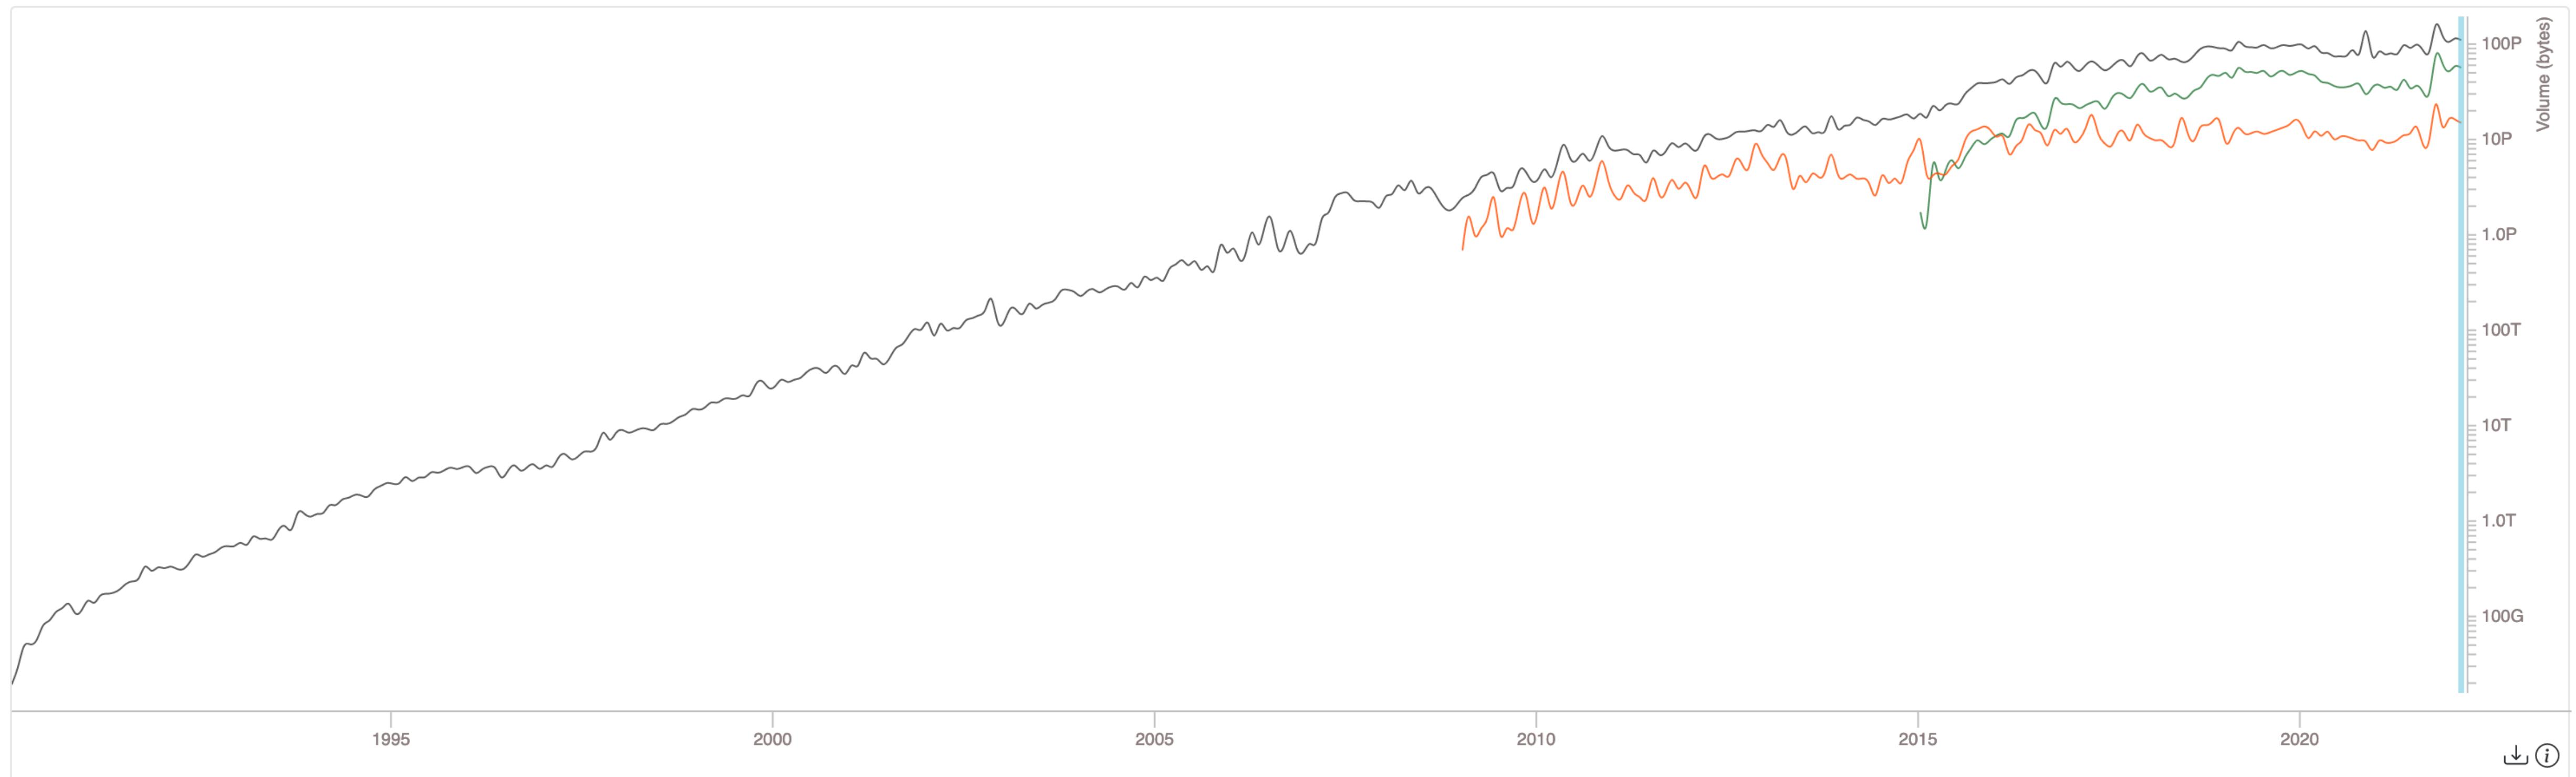
25 FEET 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21
1-10 1-11 1-12 1-13 1-14 1-15 1-16 1-17 1-18 1-19 1-20 1-21
1-22 1-23 1-24 1-25 1-26 1-27 1-28 1-29
1-33 1-34 1-35 1-36 1-37 1-38 1-39 1-40 1-41
1-42 1-43 1-44 1-45 1-46 1-47 1-48 1-49 1-50
1-51 1-52 1-53 1-54 1-55 1-56 1-57 1-58 1-59
1-60 1-61 1-62 1-63 1-64 1-65 1-66 1-67 1-68
1-69 1-70 1-71 1-72 1-73 1-74 1-75 1-76 1-77
1-78 1-79 1-80 1-81 1-82 1-83 1-84 1-85 1-86
1-87 1-88 1-89 1-90 1-91 1-92 1-93 1-94 1-95
1-96 1-97 1-98 1-99 1-100

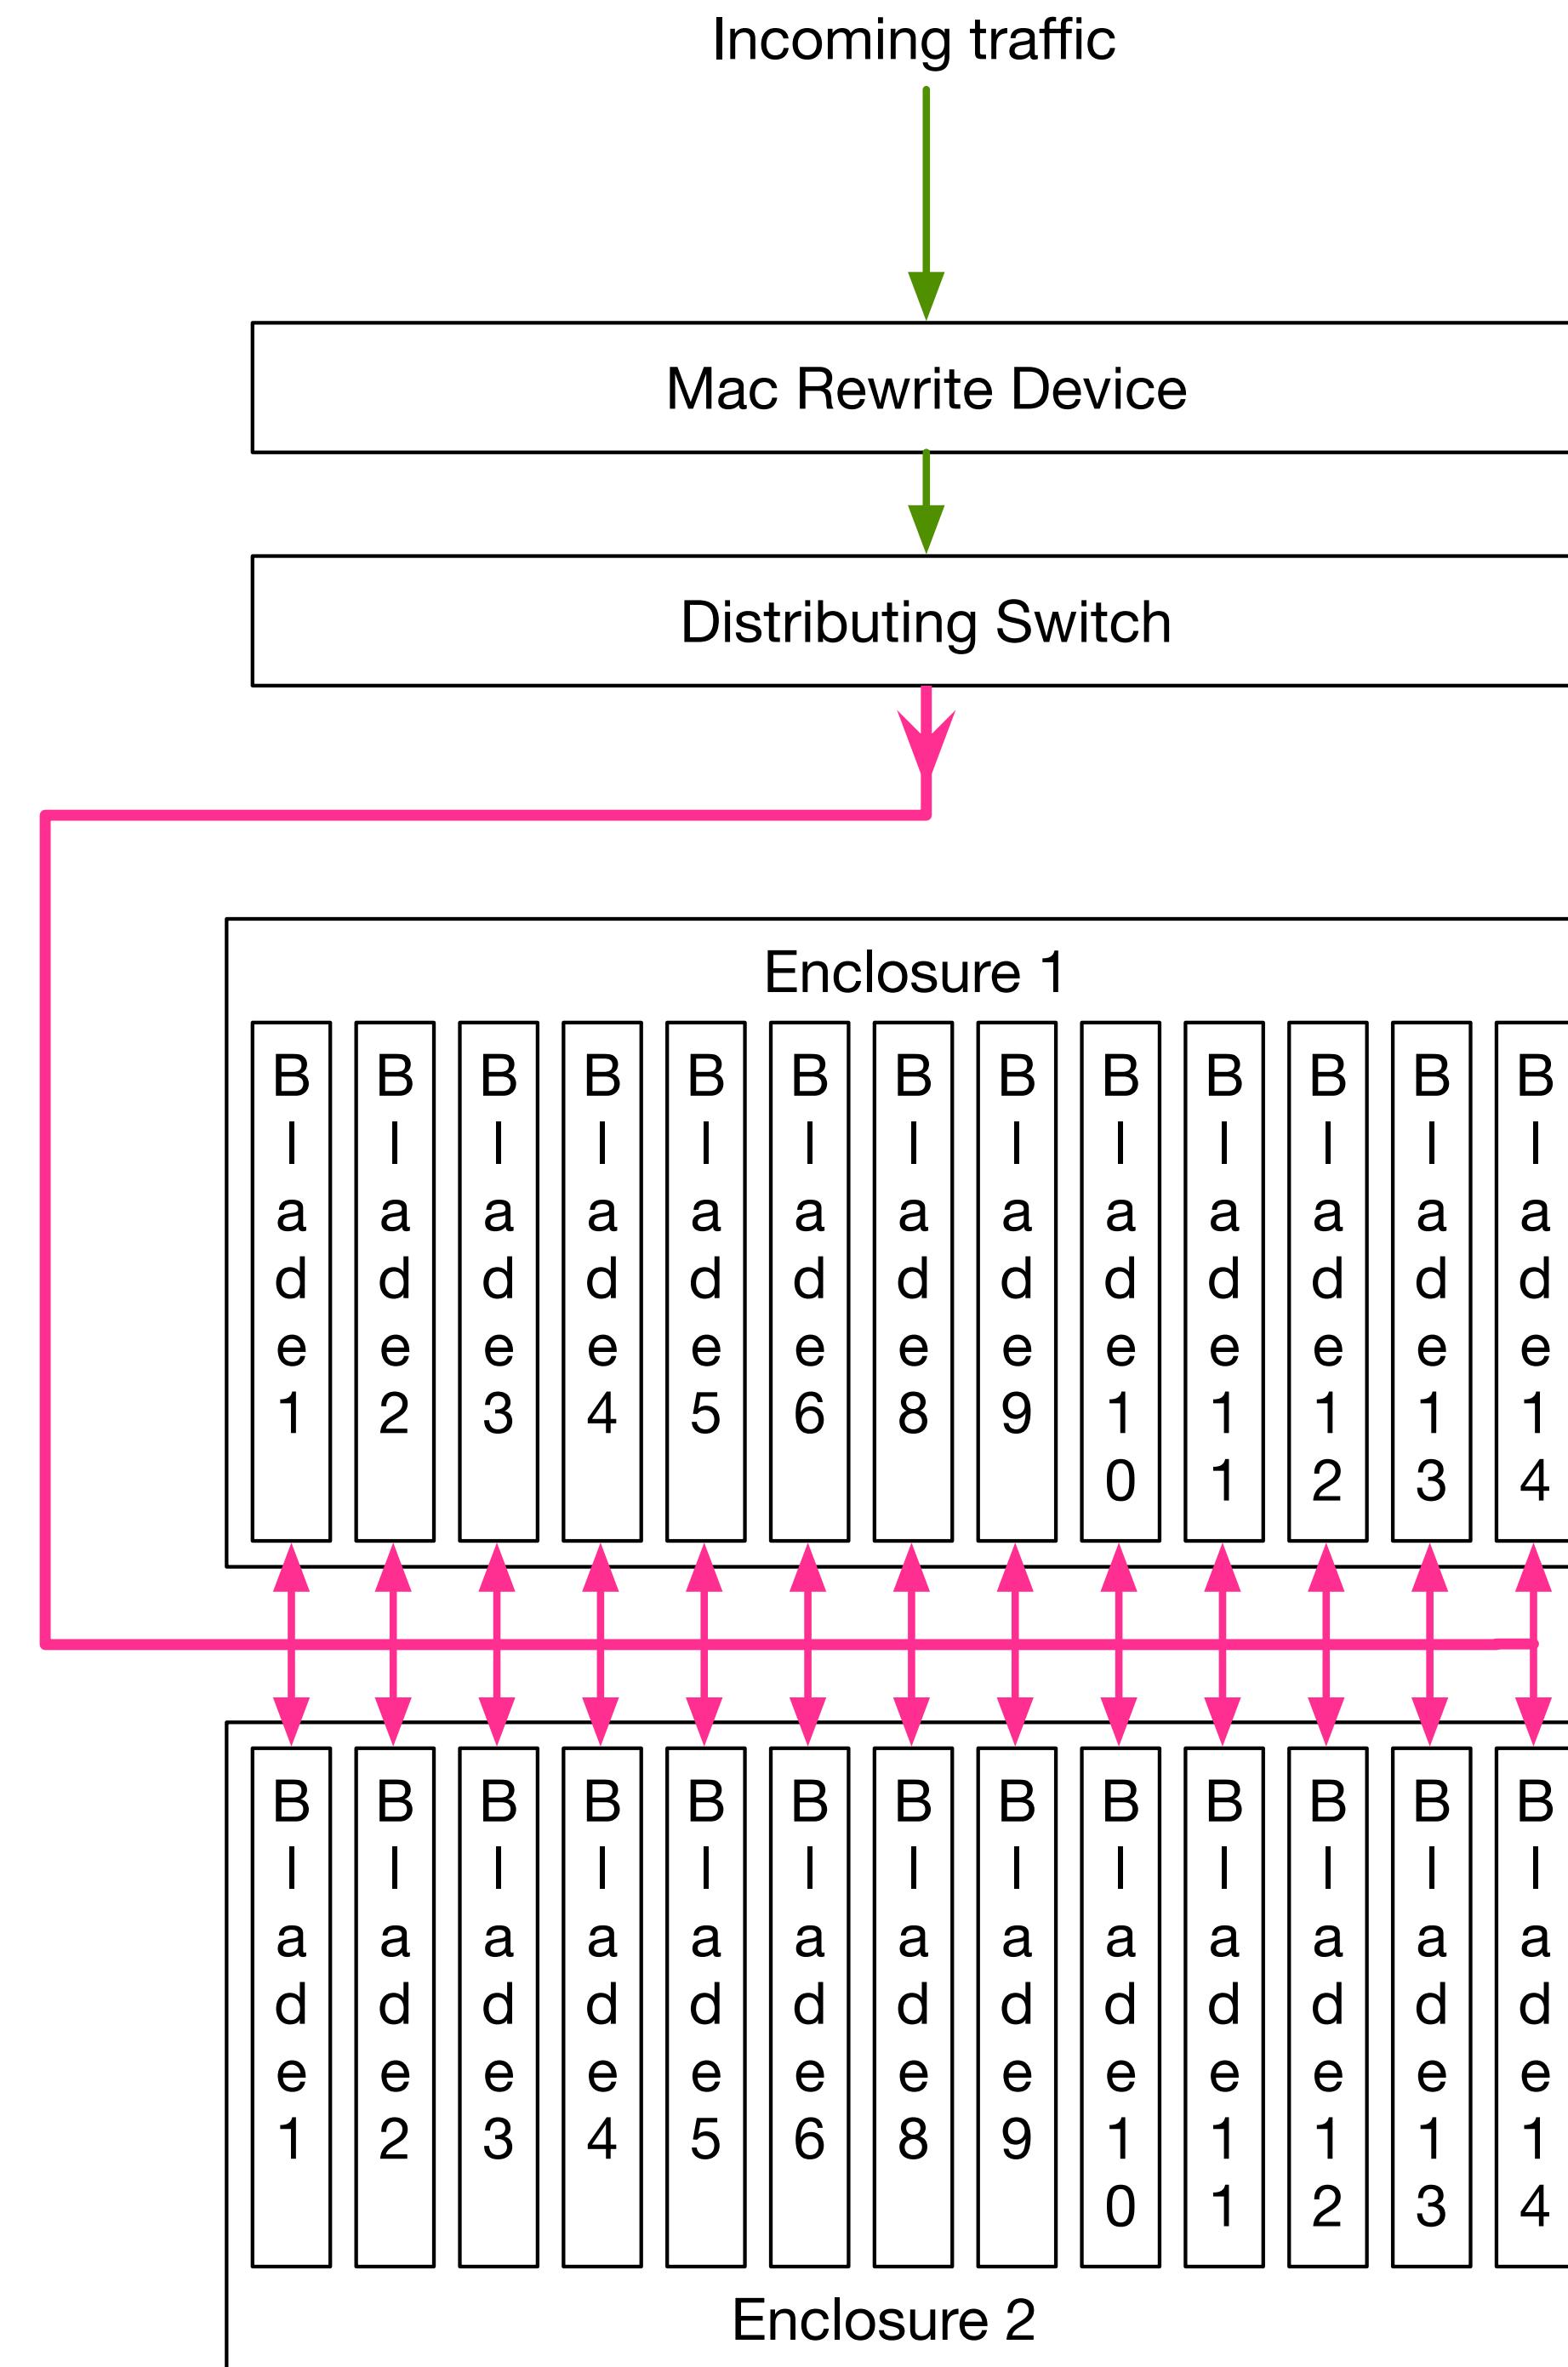
25 FEET 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21
1-10 1-11 1-12 1-13 1-14 1-15 1-16 1-17 1-18 1-19 1-20 1-21
1-22 1-23 1-24 1-25 1-26 1-27 1-28 1-29
1-33 1-34 1-35 1-36 1-37 1-38 1-39 1-40 1-41
1-42 1-43 1-44 1-45 1-46 1-47 1-48 1-49 1-50
1-51 1-52 1-53 1-54 1-55 1-56 1-57 1-58 1-59
1-60 1-61 1-62 1-63 1-64 1-65 1-66 1-67 1-68
1-69 1-70 1-71 1-72 1-73 1-74 1-75 1-76 1-77
1-78 1-79 1-80 1-81 1-82 1-83 1-84 1-85 1-86
1-87 1-88 1-89 1-90 1-91 1-92 1-93 1-94 1-95
1-96 1-97 1-98 1-99 1-100

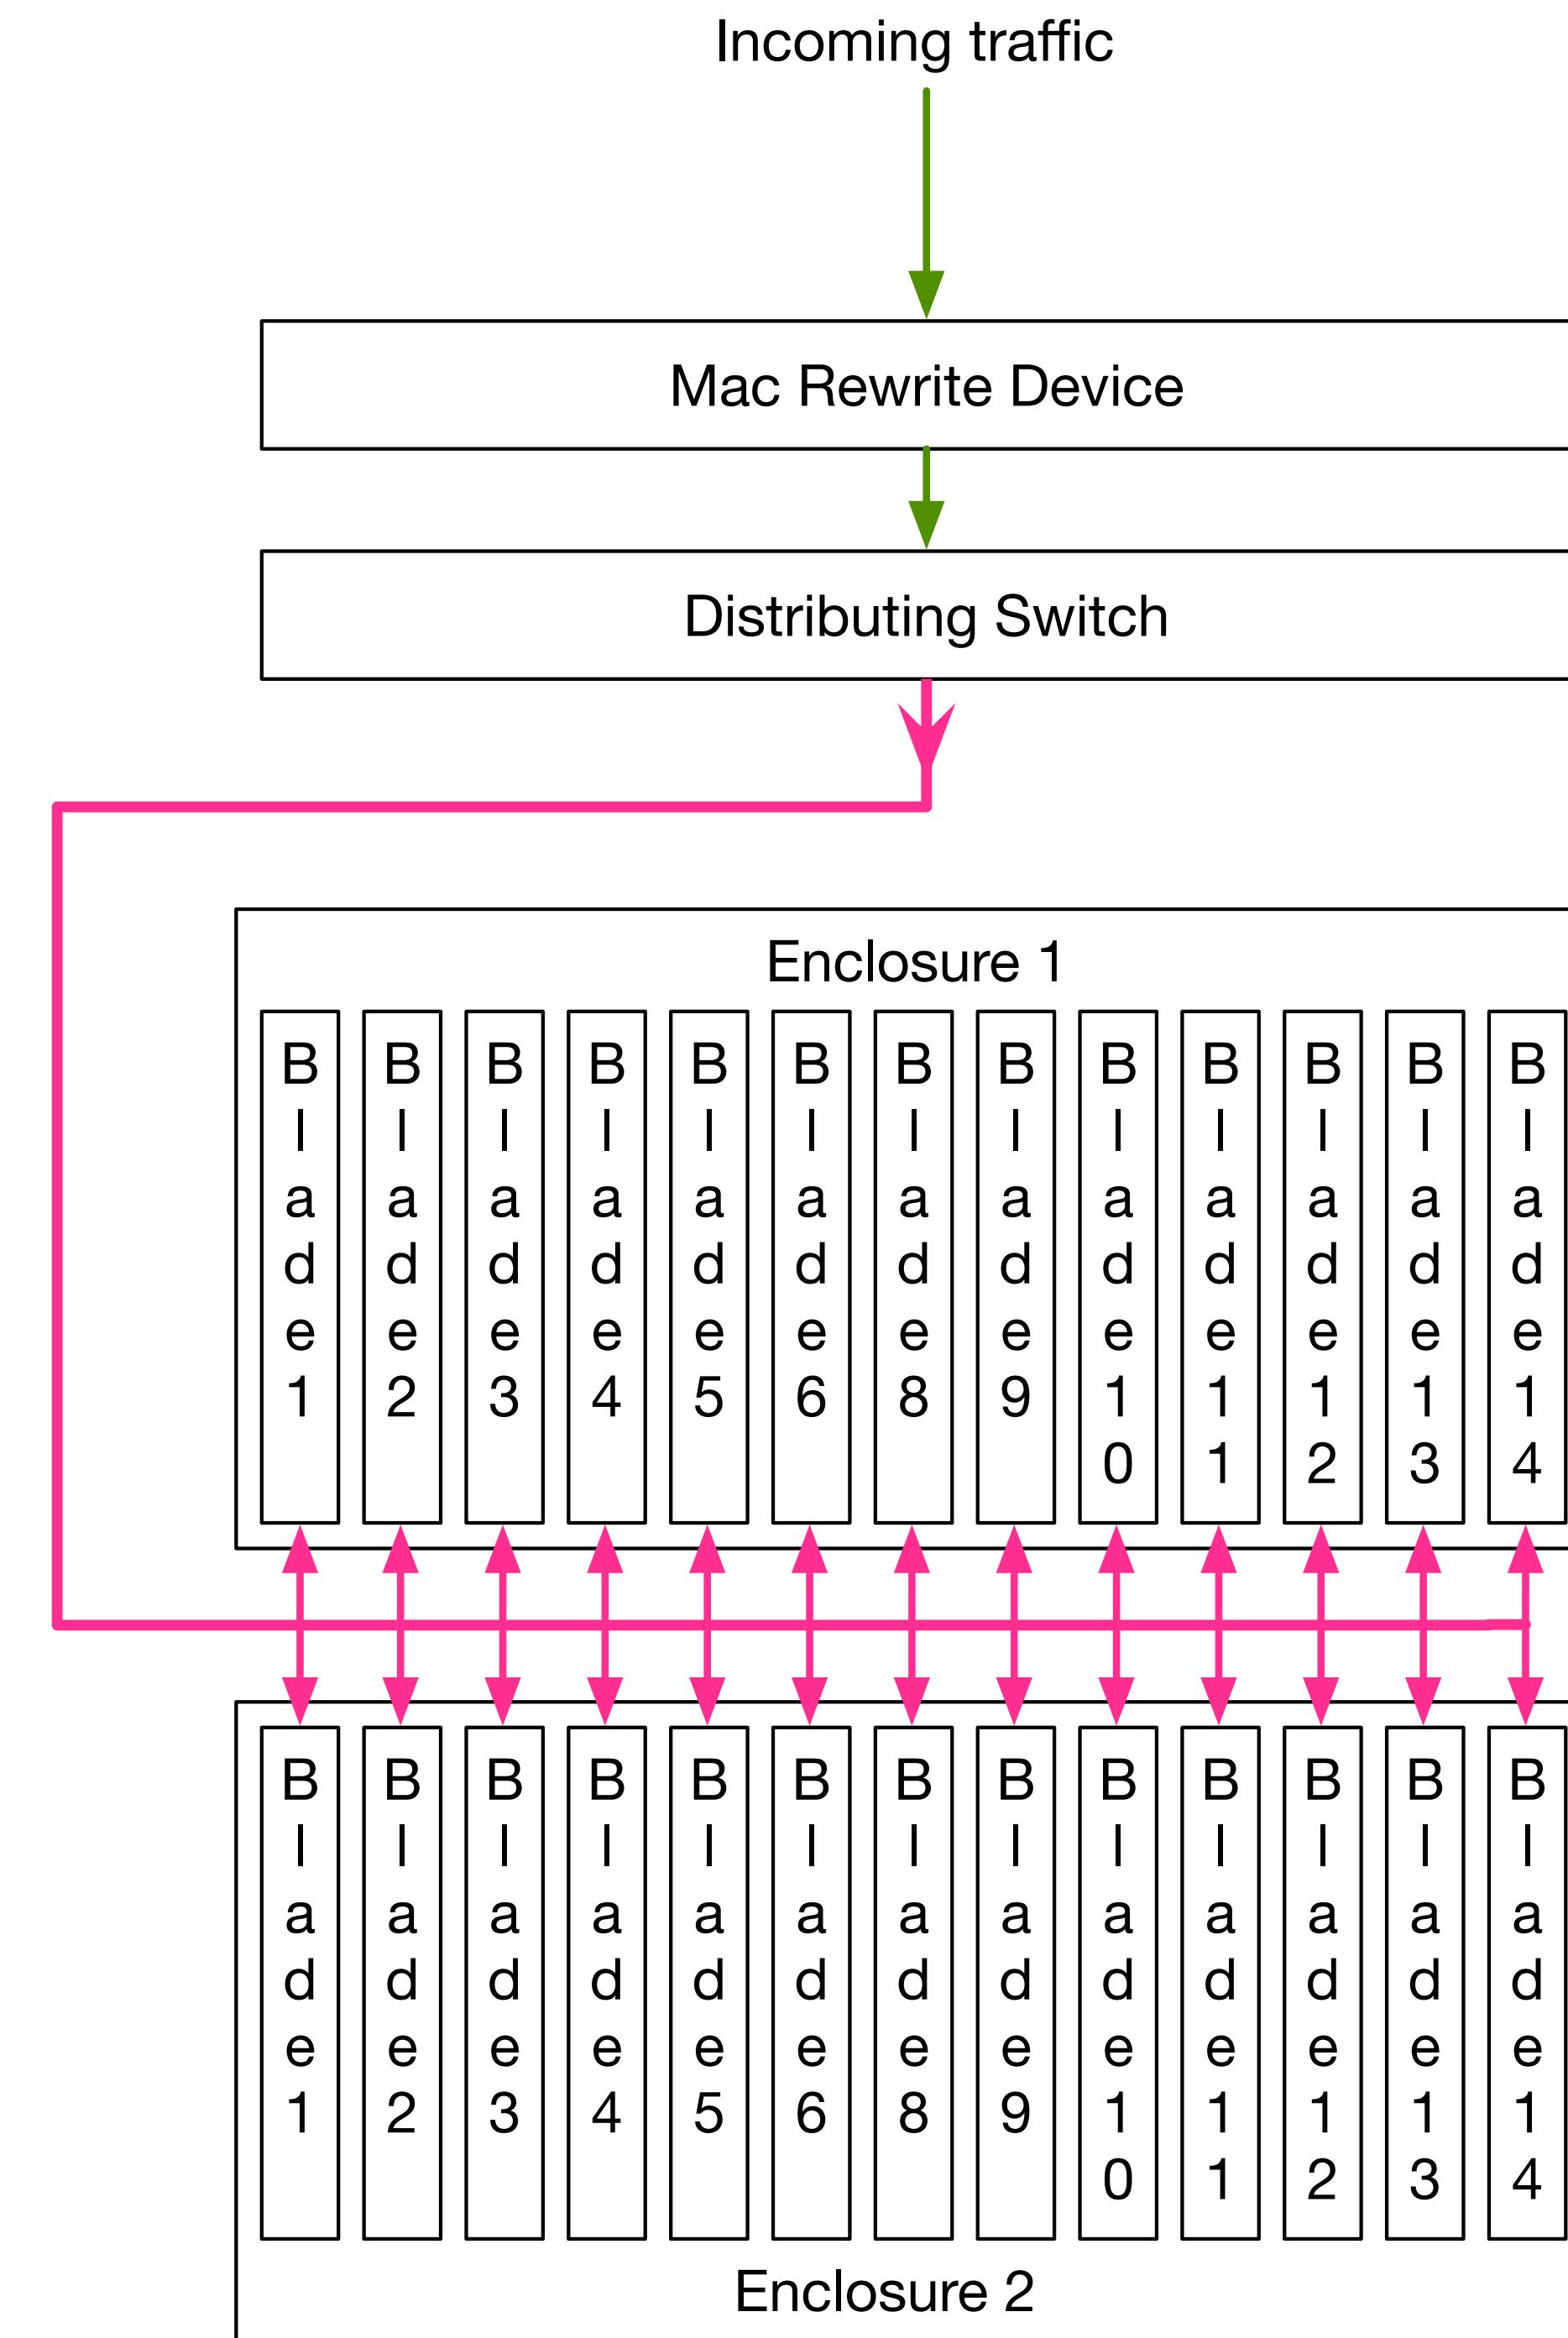
Typical Network Monitoring Setup

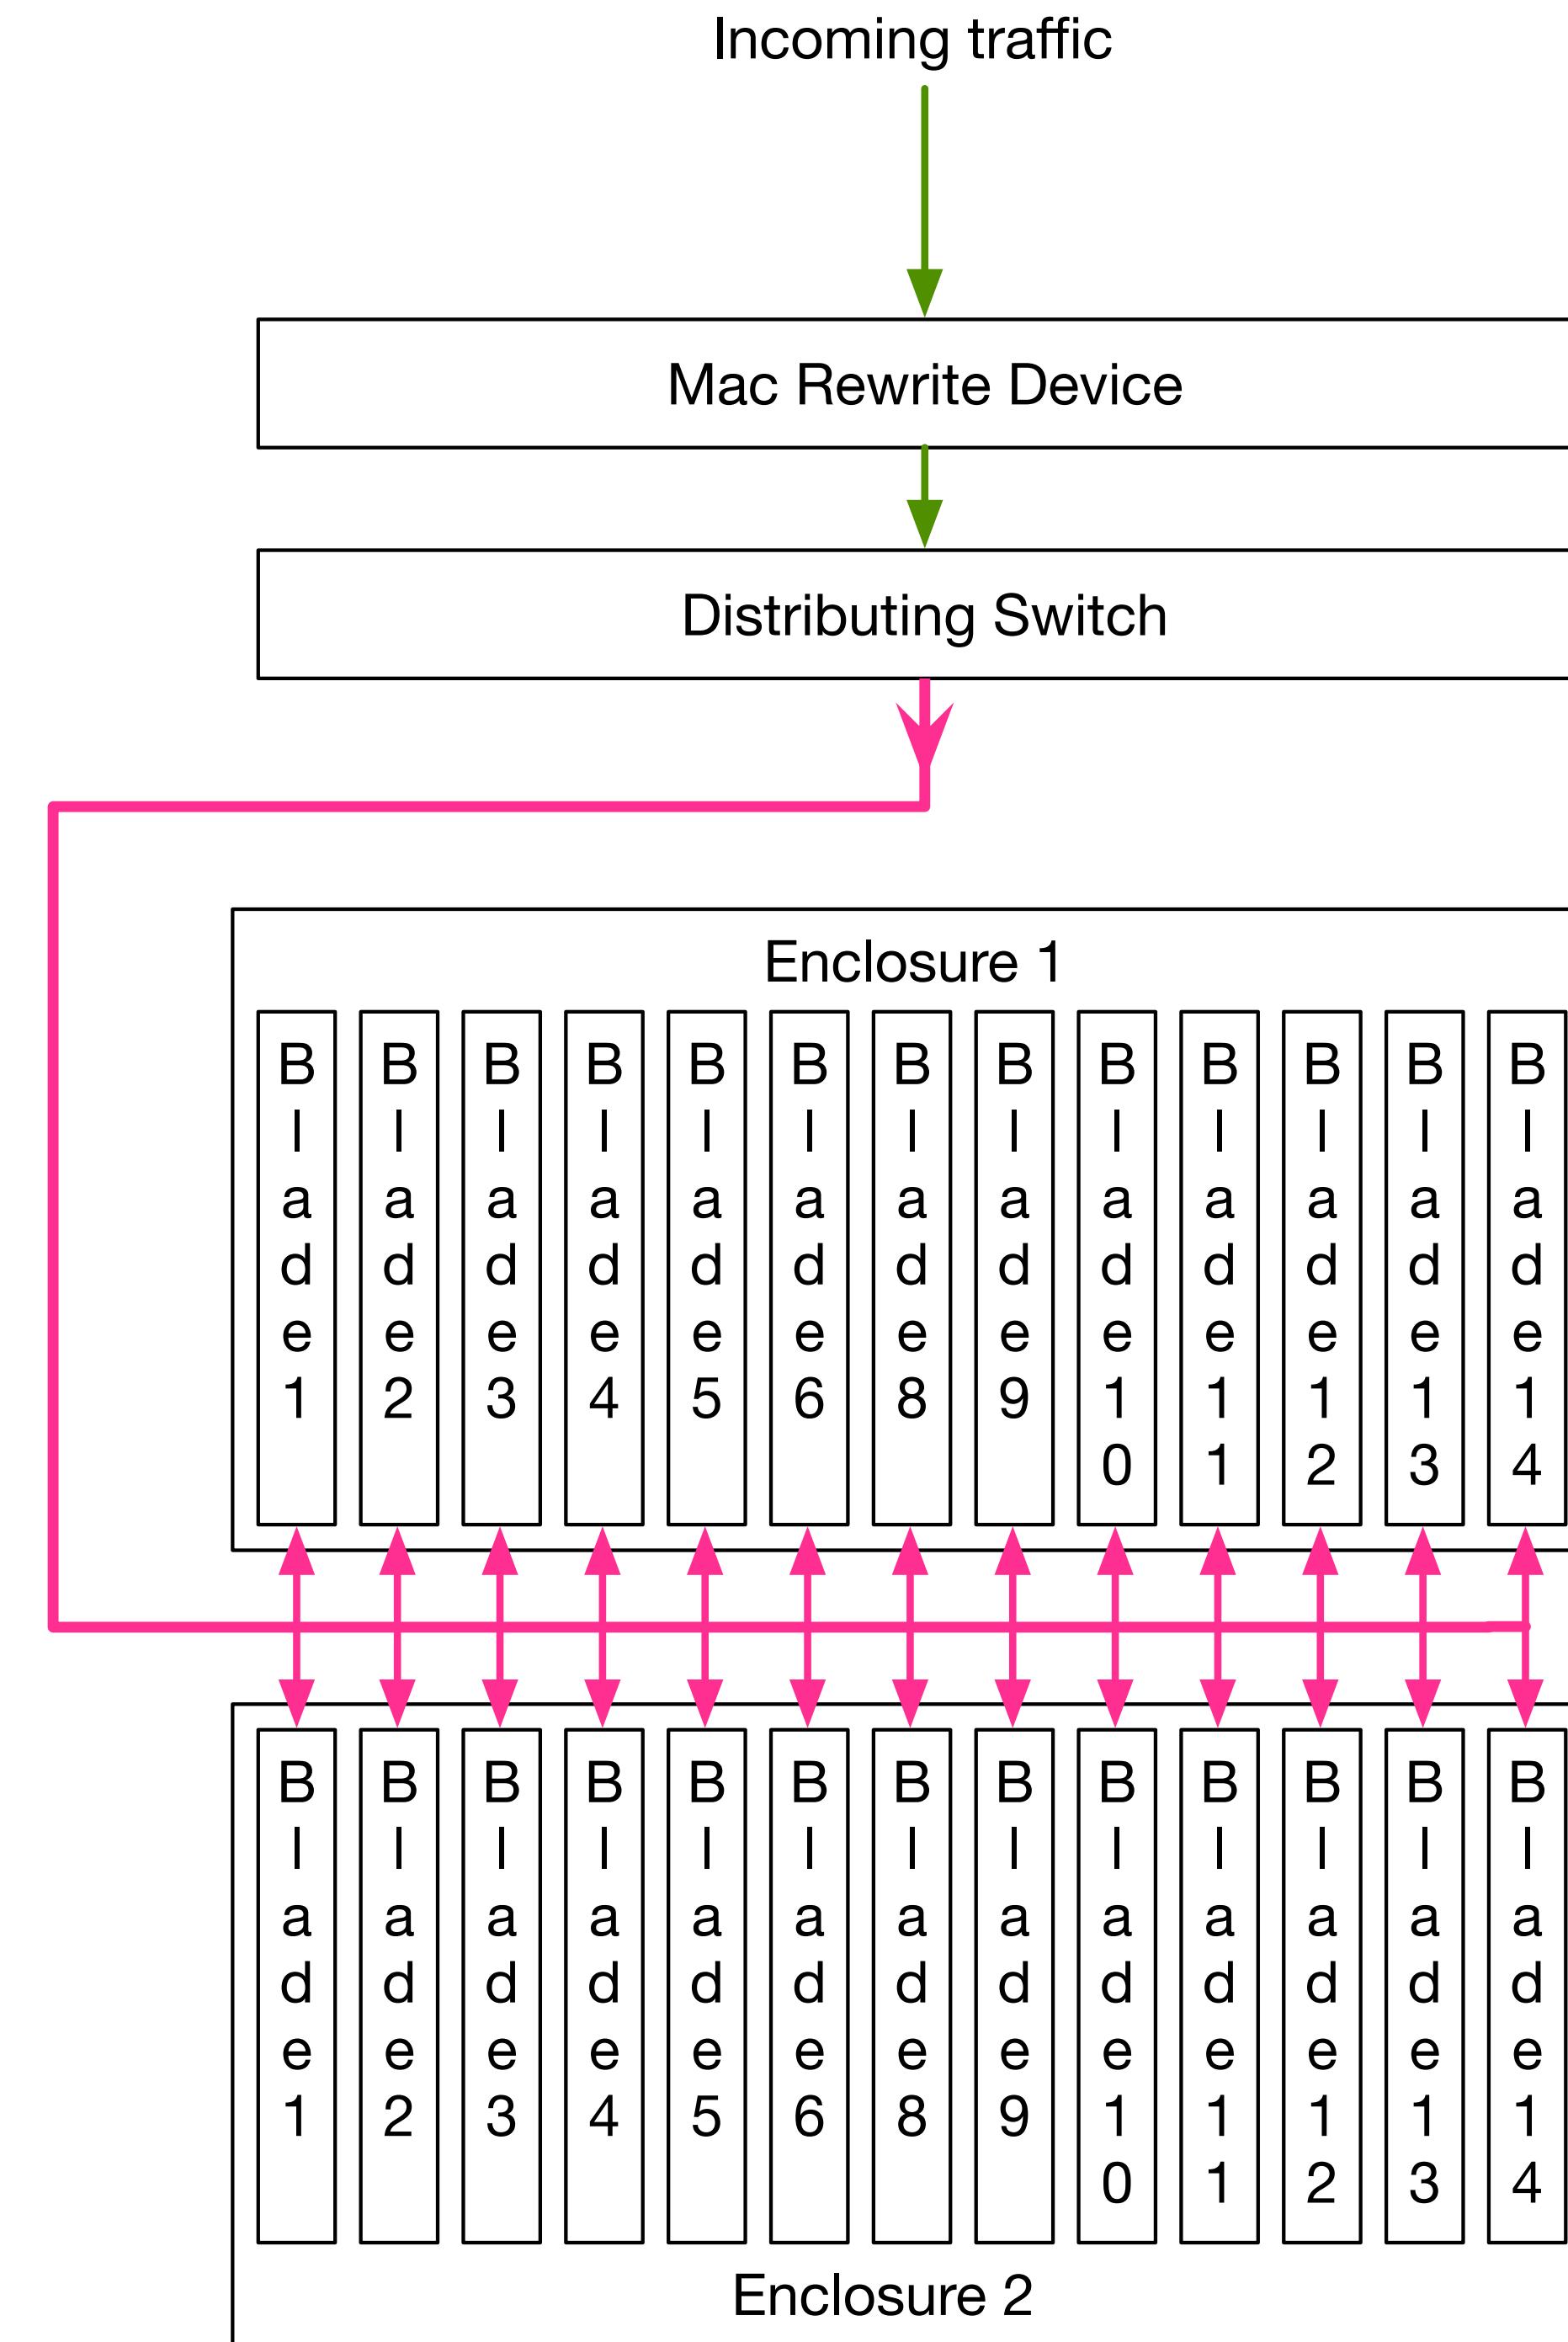


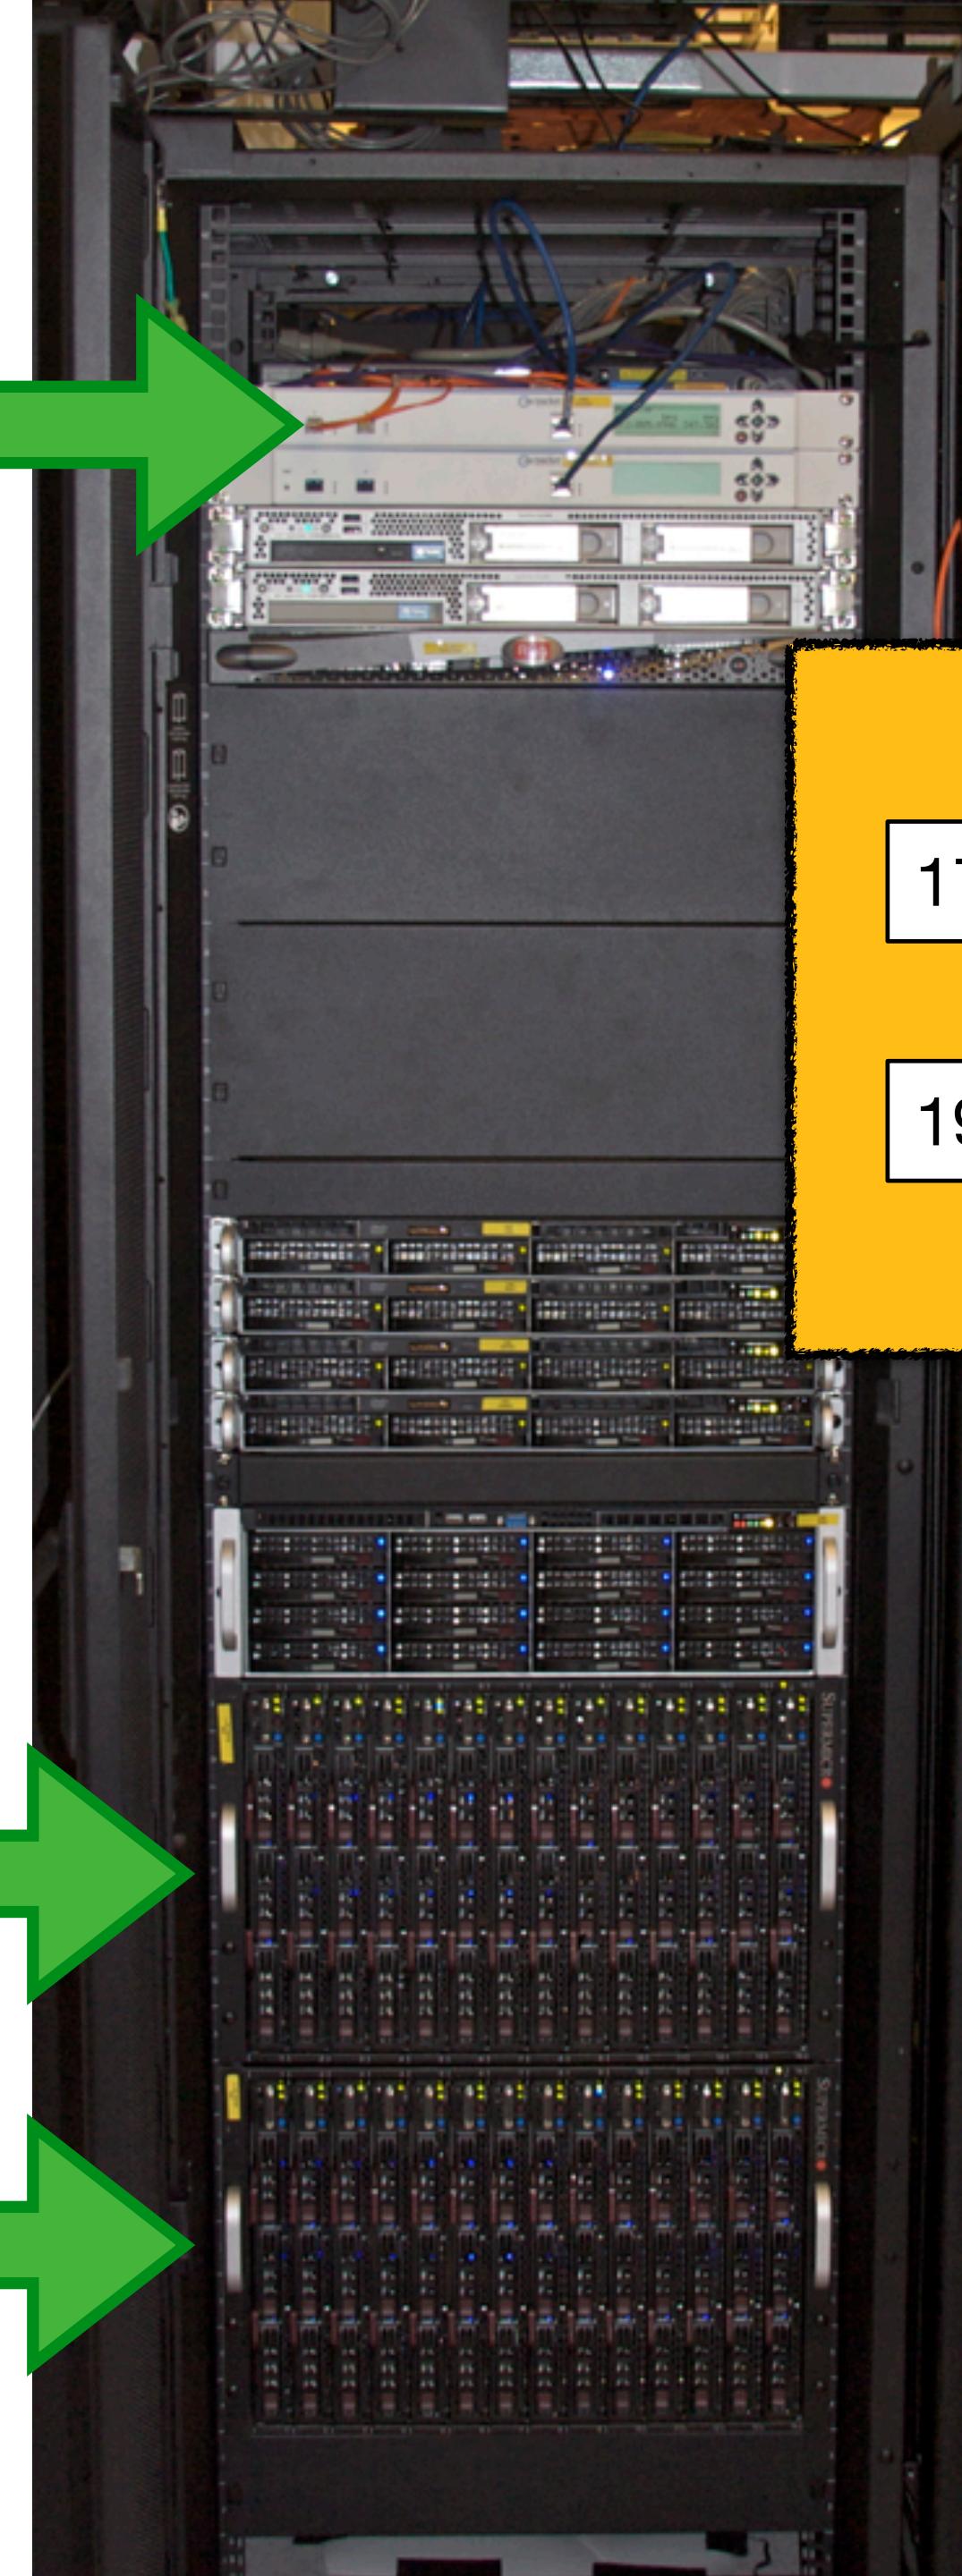












Incoming traffic

Mac Rewrite Device

172.29.12.12 | 37862 | 192.0.78.212 | 443 | TCP

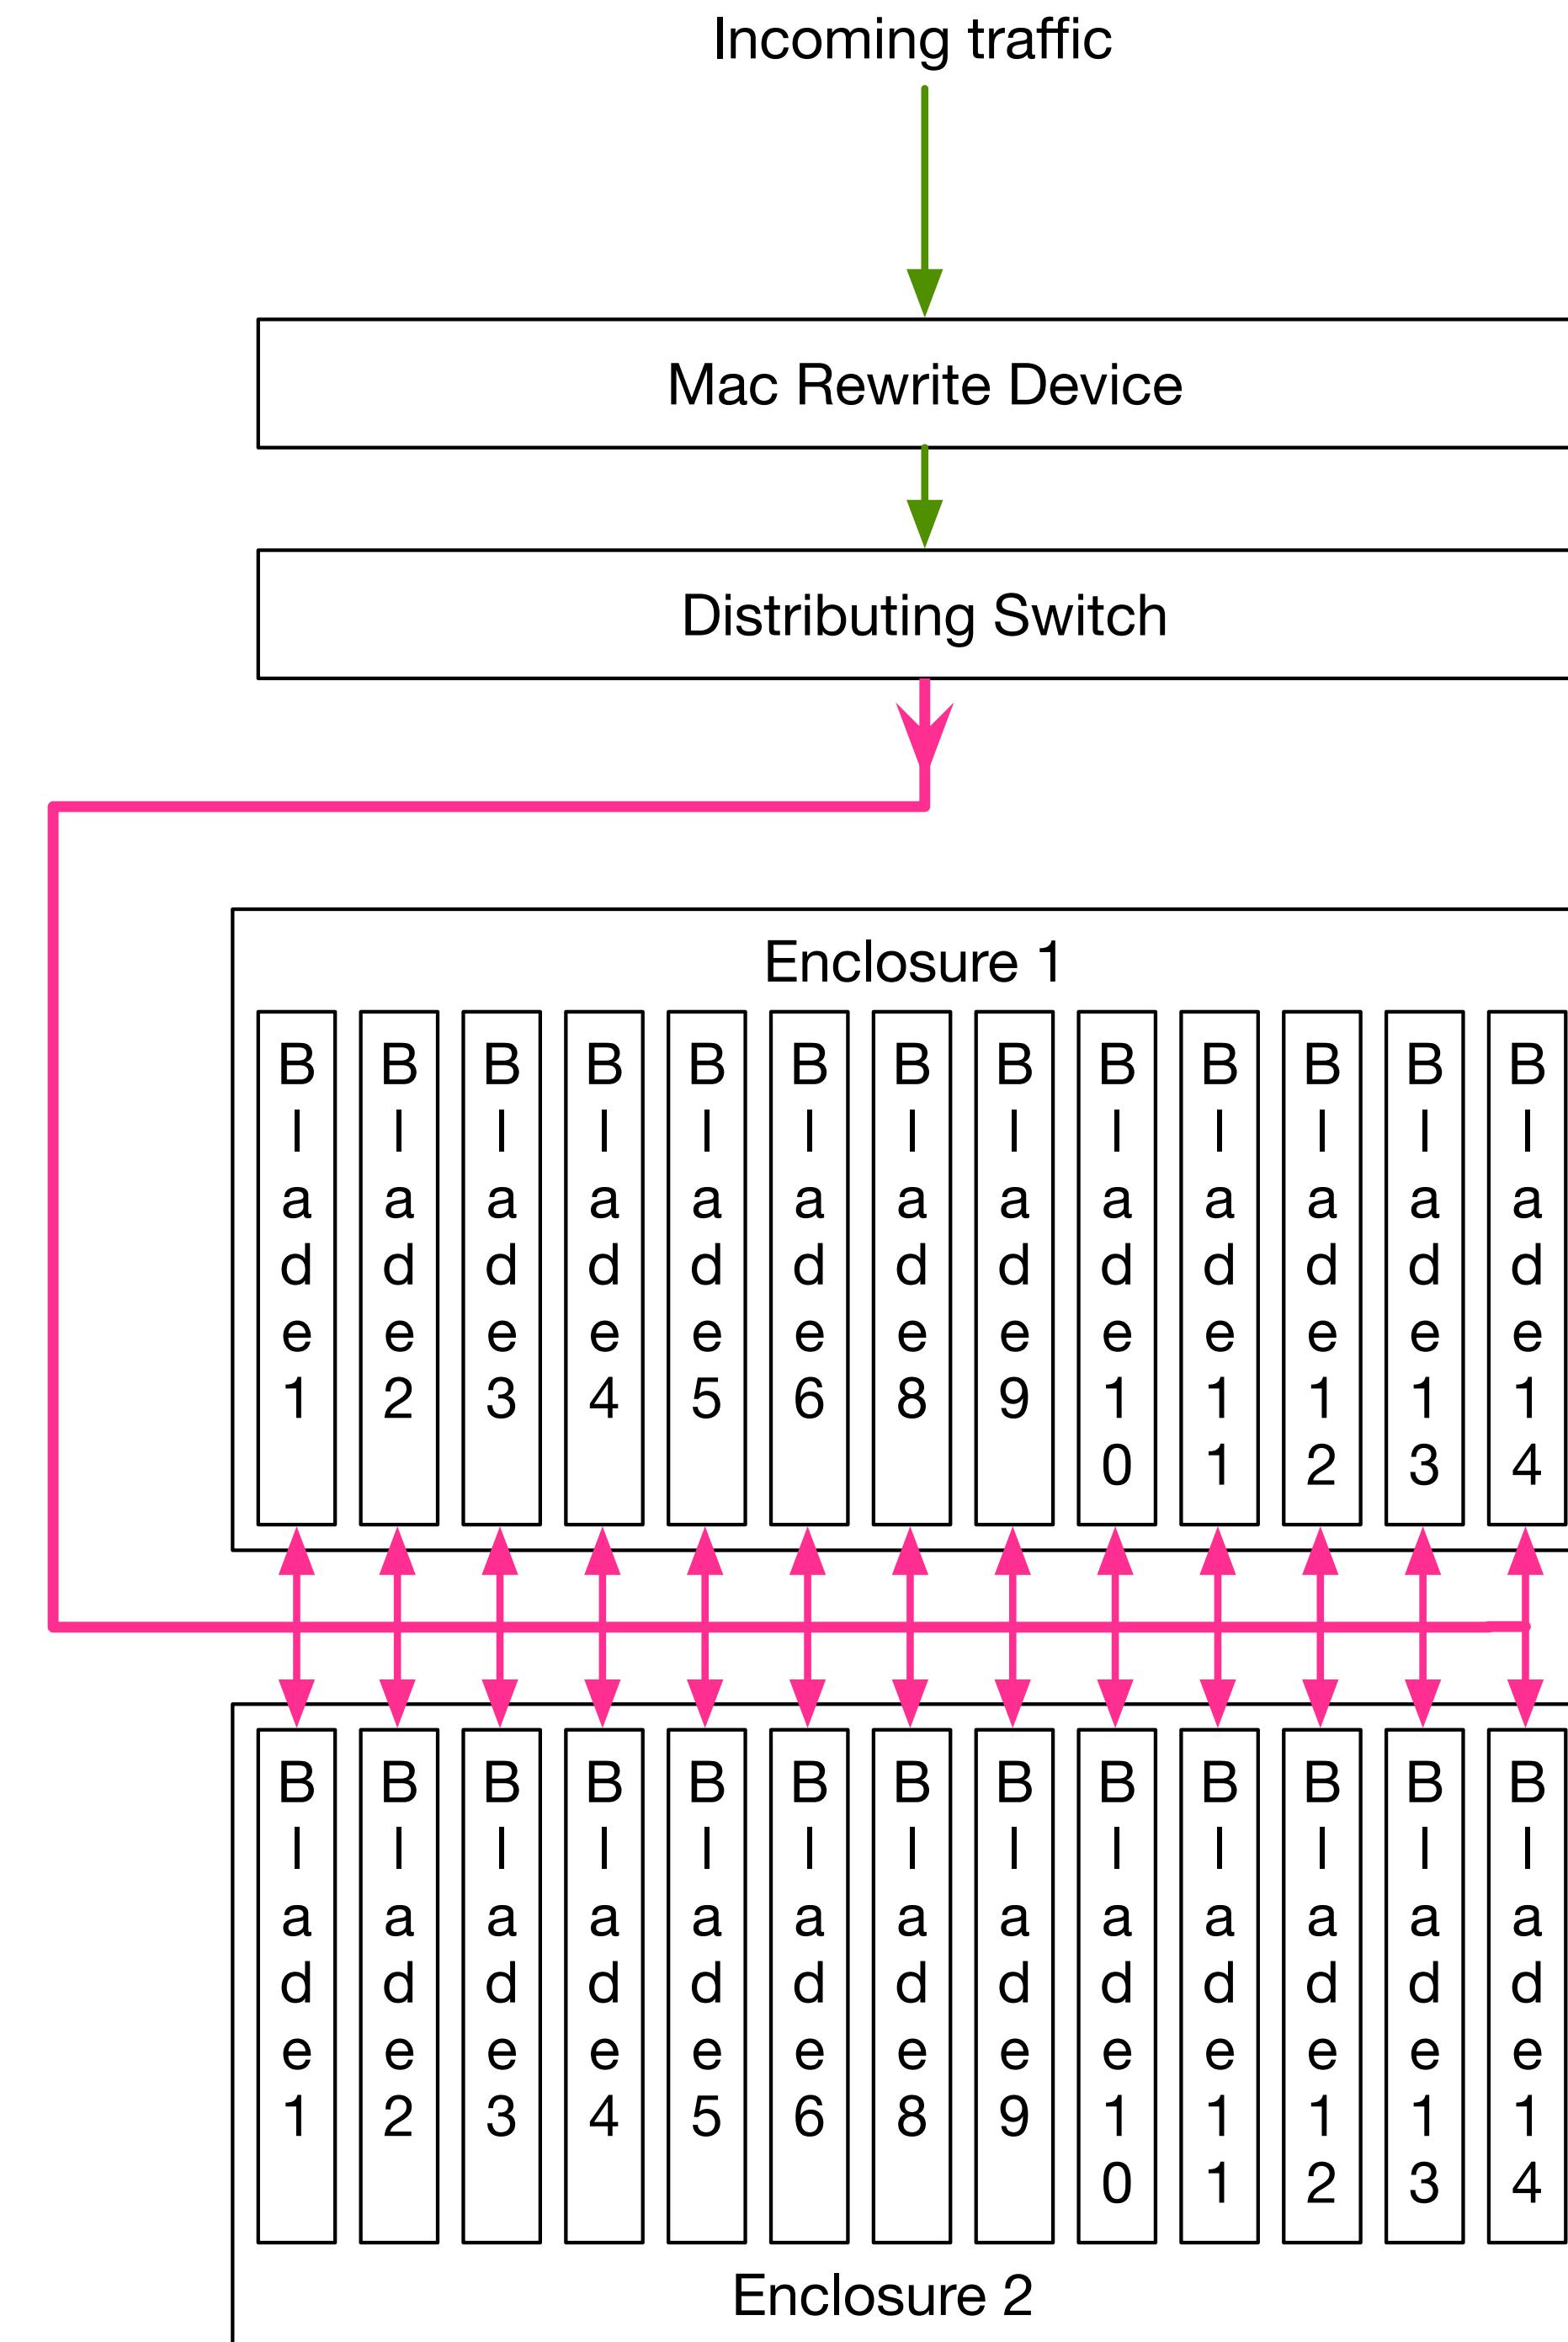
192.0.78.212 | 443 | 172.29.12.12 | 37862 | TCP

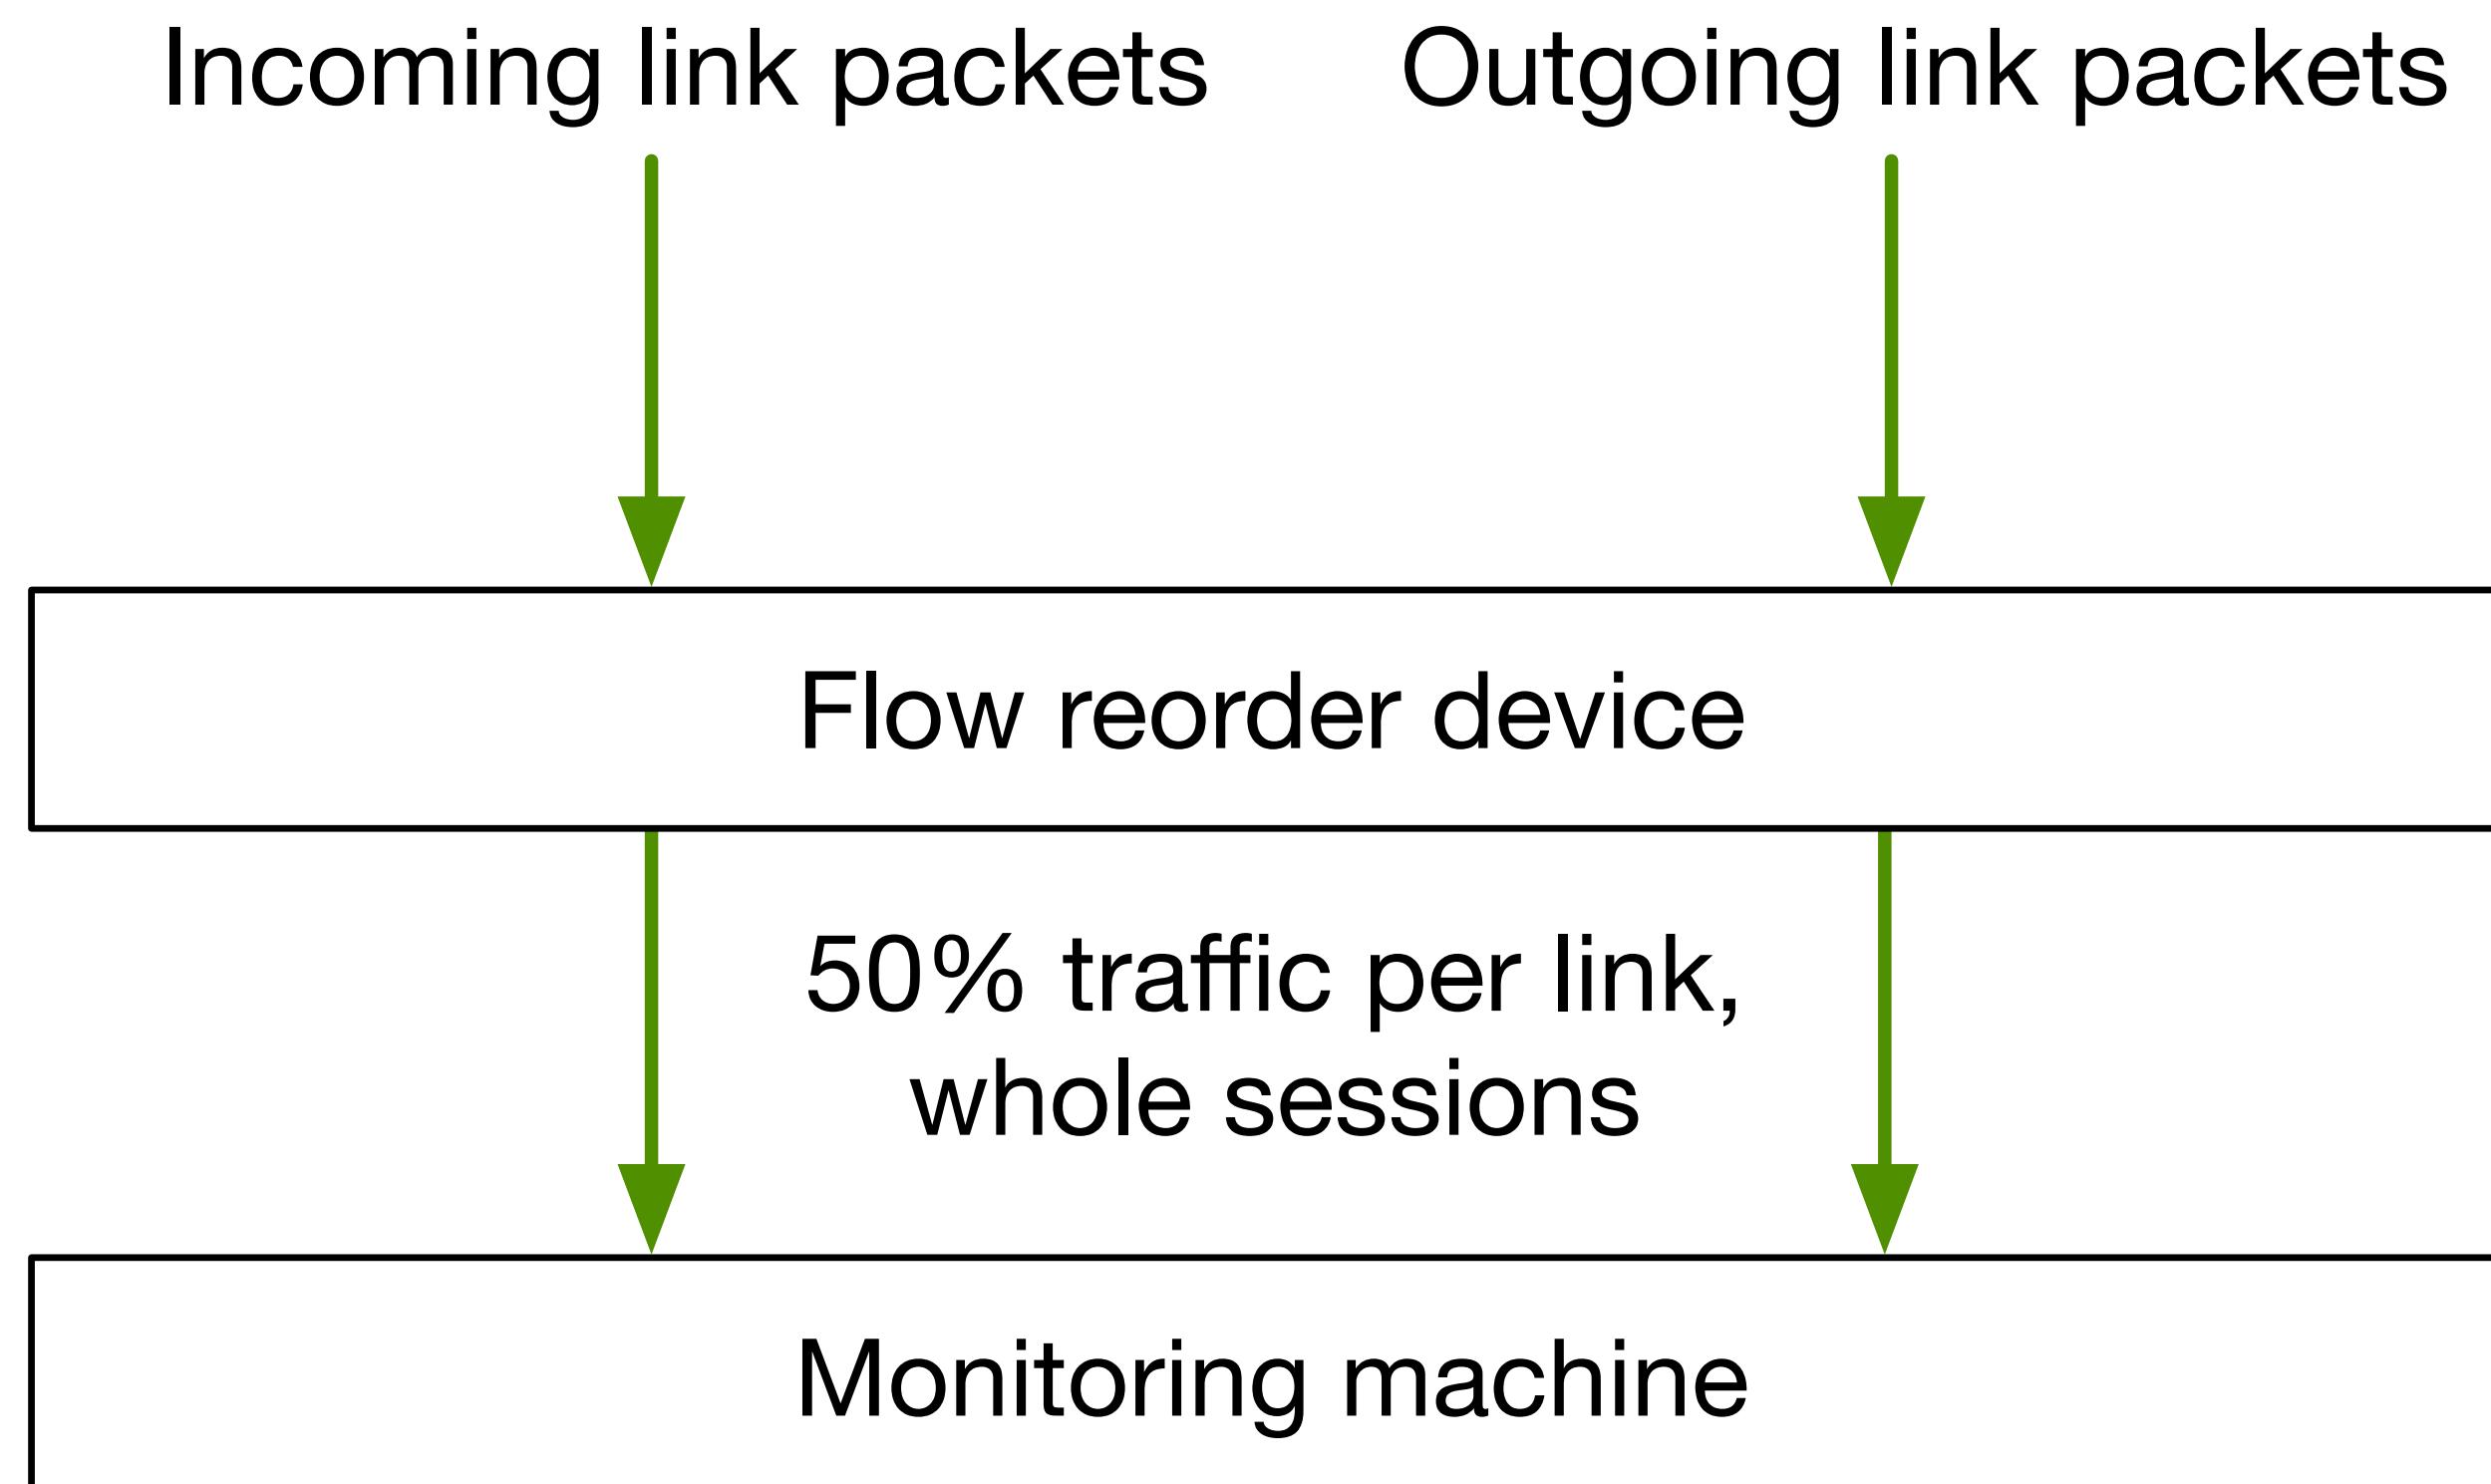
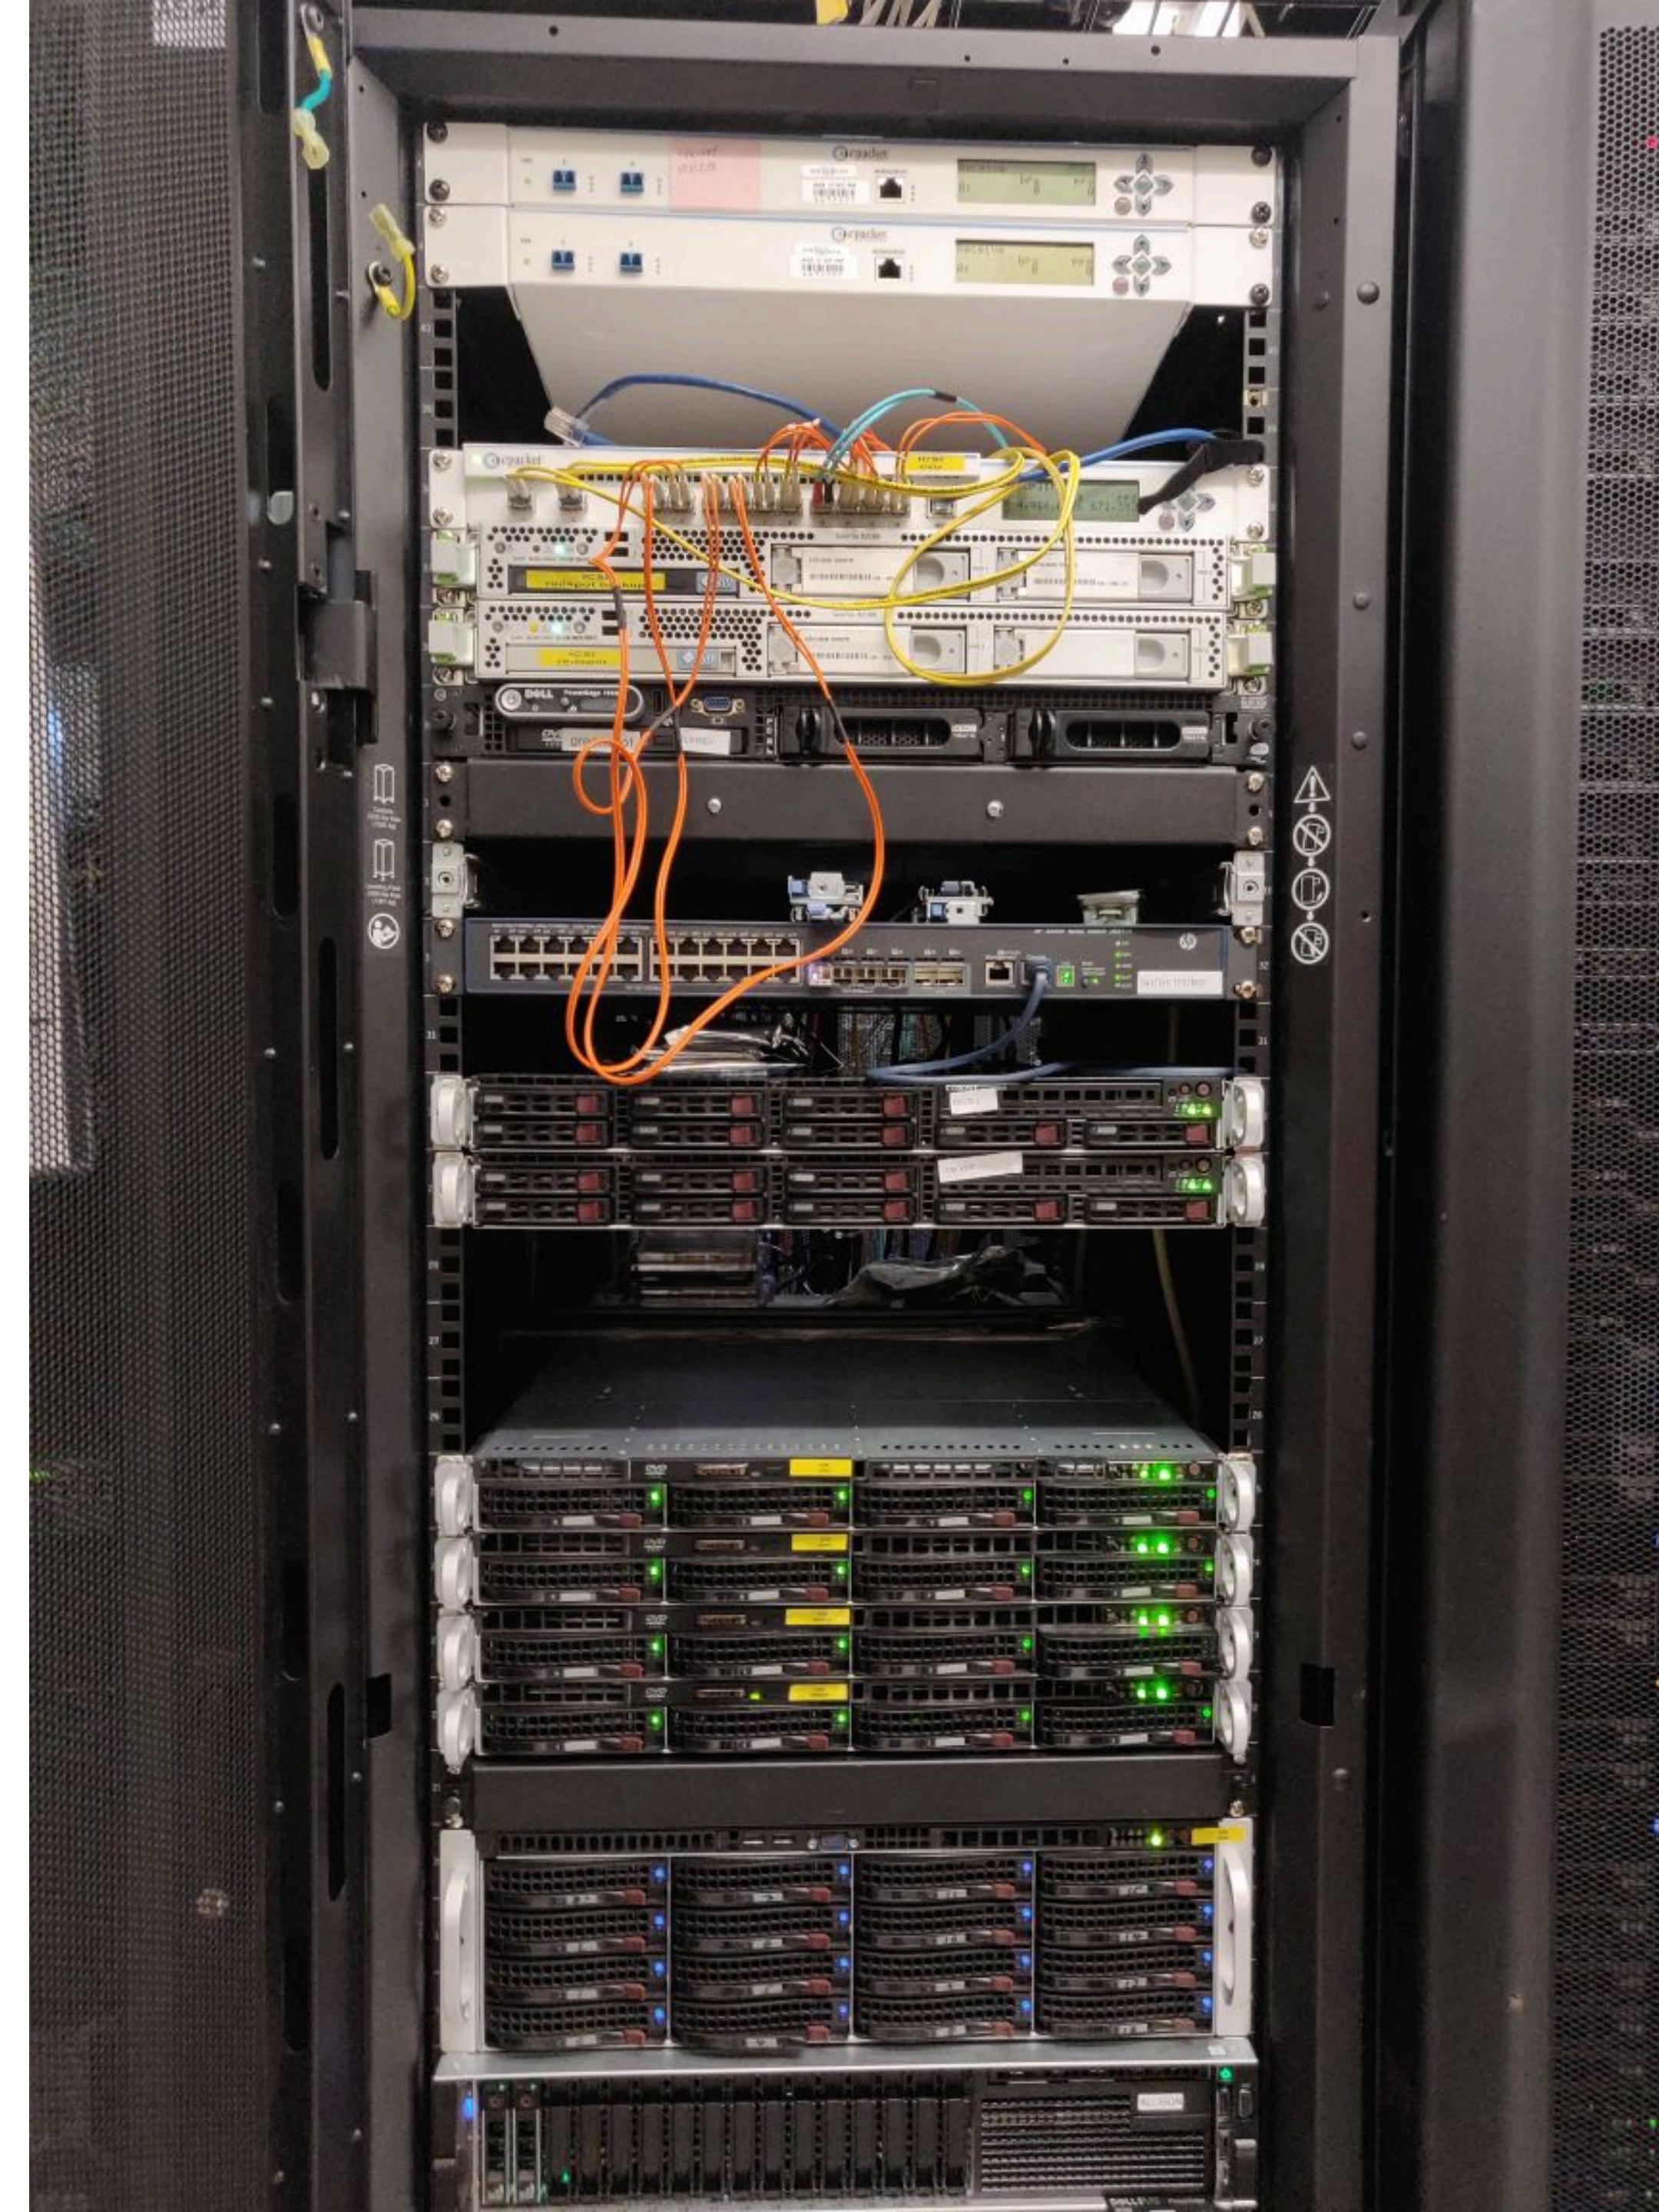
3fa795...

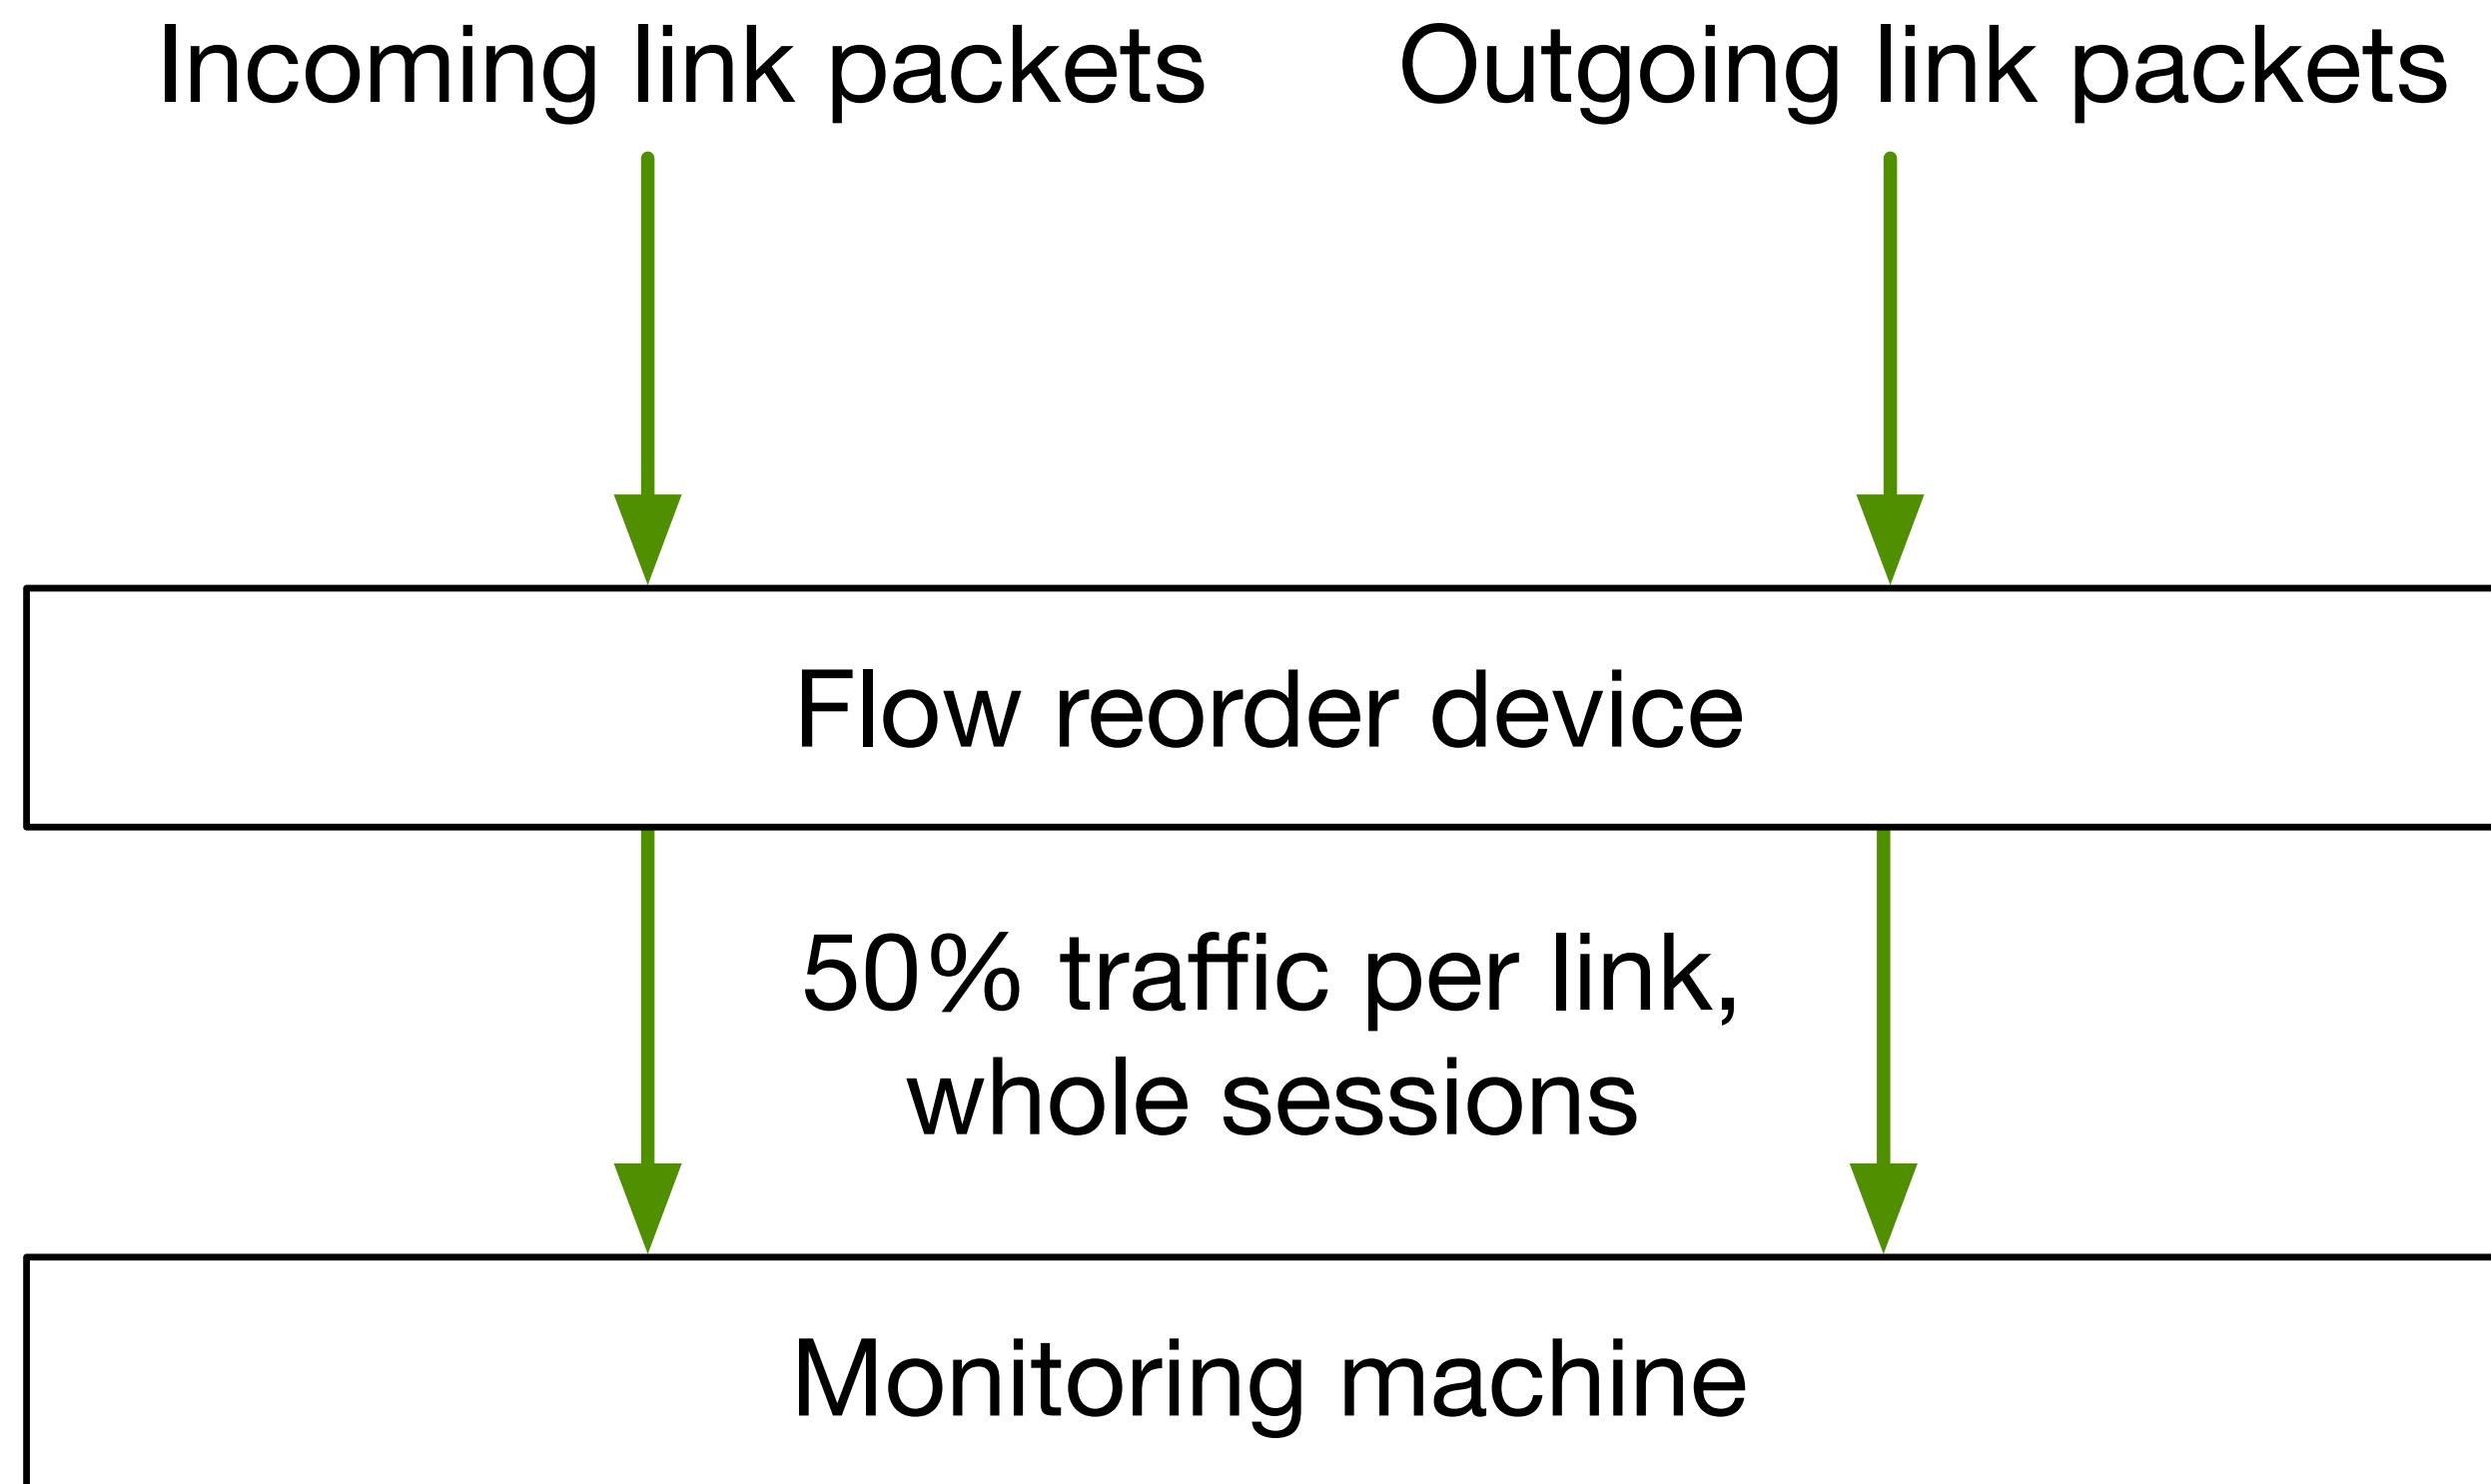
d
e
1
d
e
2
d
e
3
d
e
4
d
e
5
d
e
6
d
e
8
d
e
9
d
e
1
0
d
e
1
1
d
e
1
2
d
e
1
3
d
e
1
4

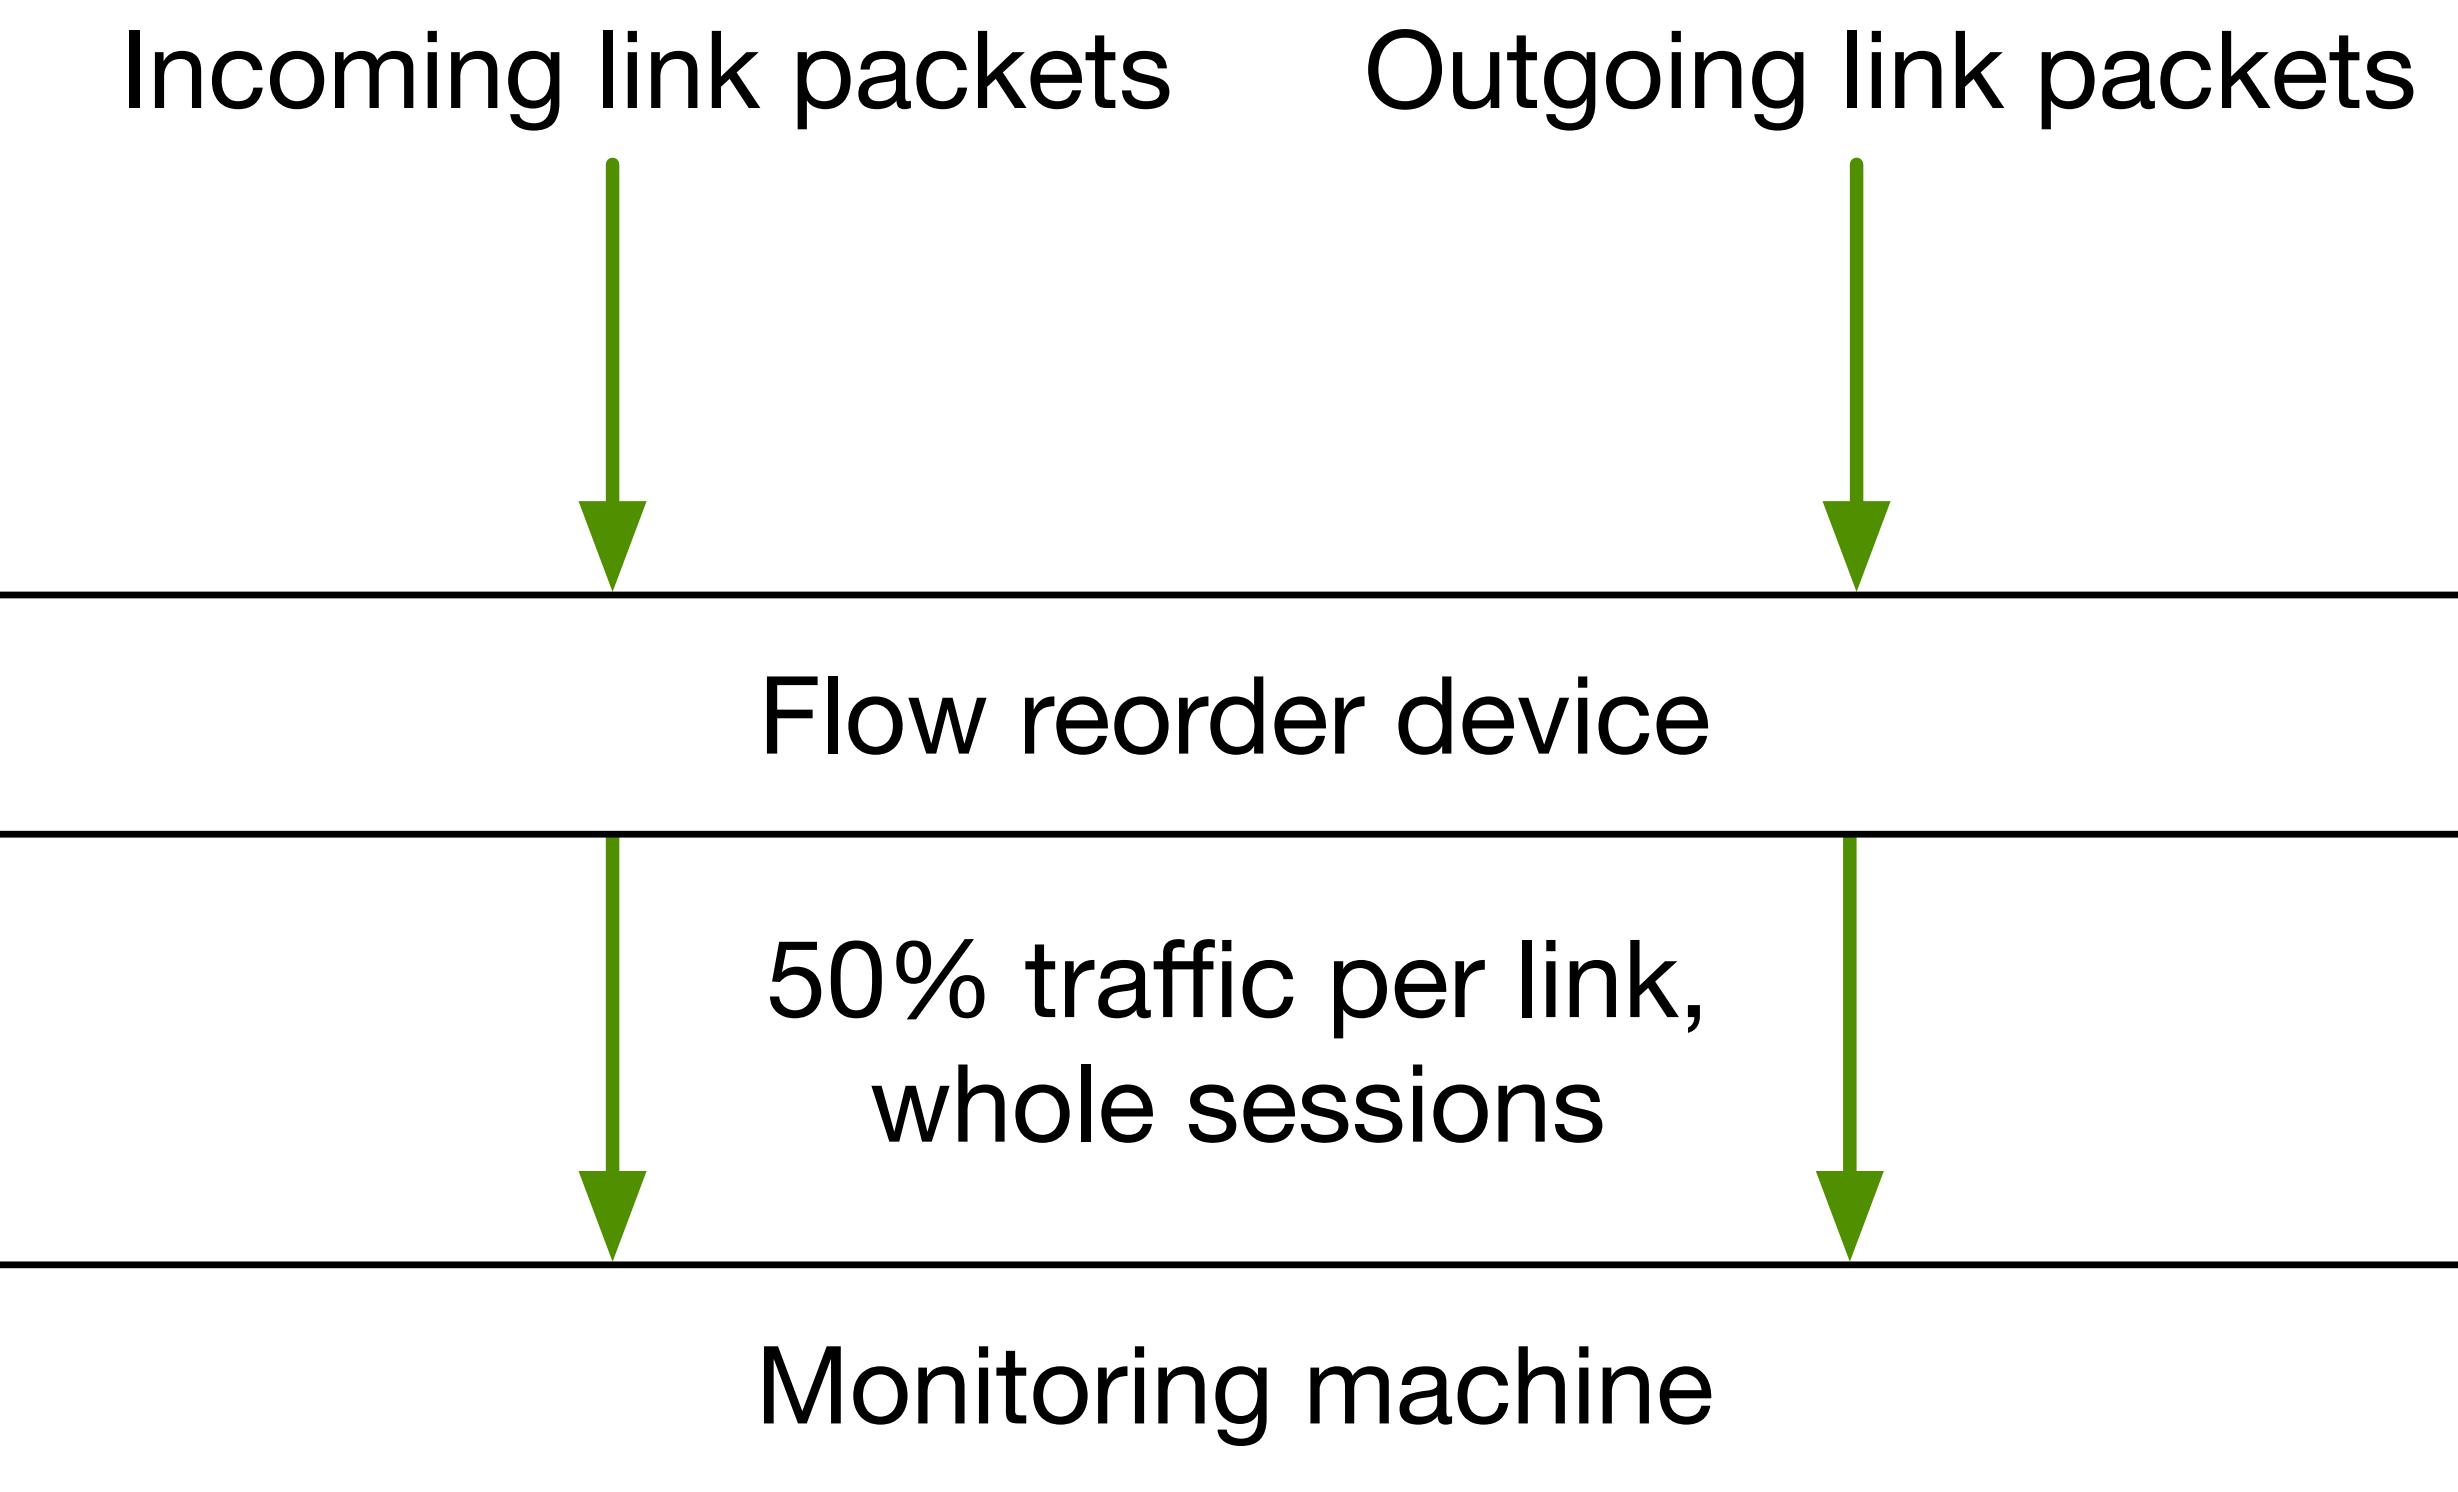
B
I
a
d
e
1
B
I
a
d
e
2
B
I
a
d
e
3
B
I
a
d
e
4
B
I
a
d
e
5
B
I
a
d
e
6
B
I
a
d
e
8
B
I
a
d
e
9
B
I
a
d
e
1
0
B
I
a
d
e
1
1
B
I
a
d
e
1
2
B
I
a
d
e
1
3
B
I
a
d
e
1
4

Enclosure 2

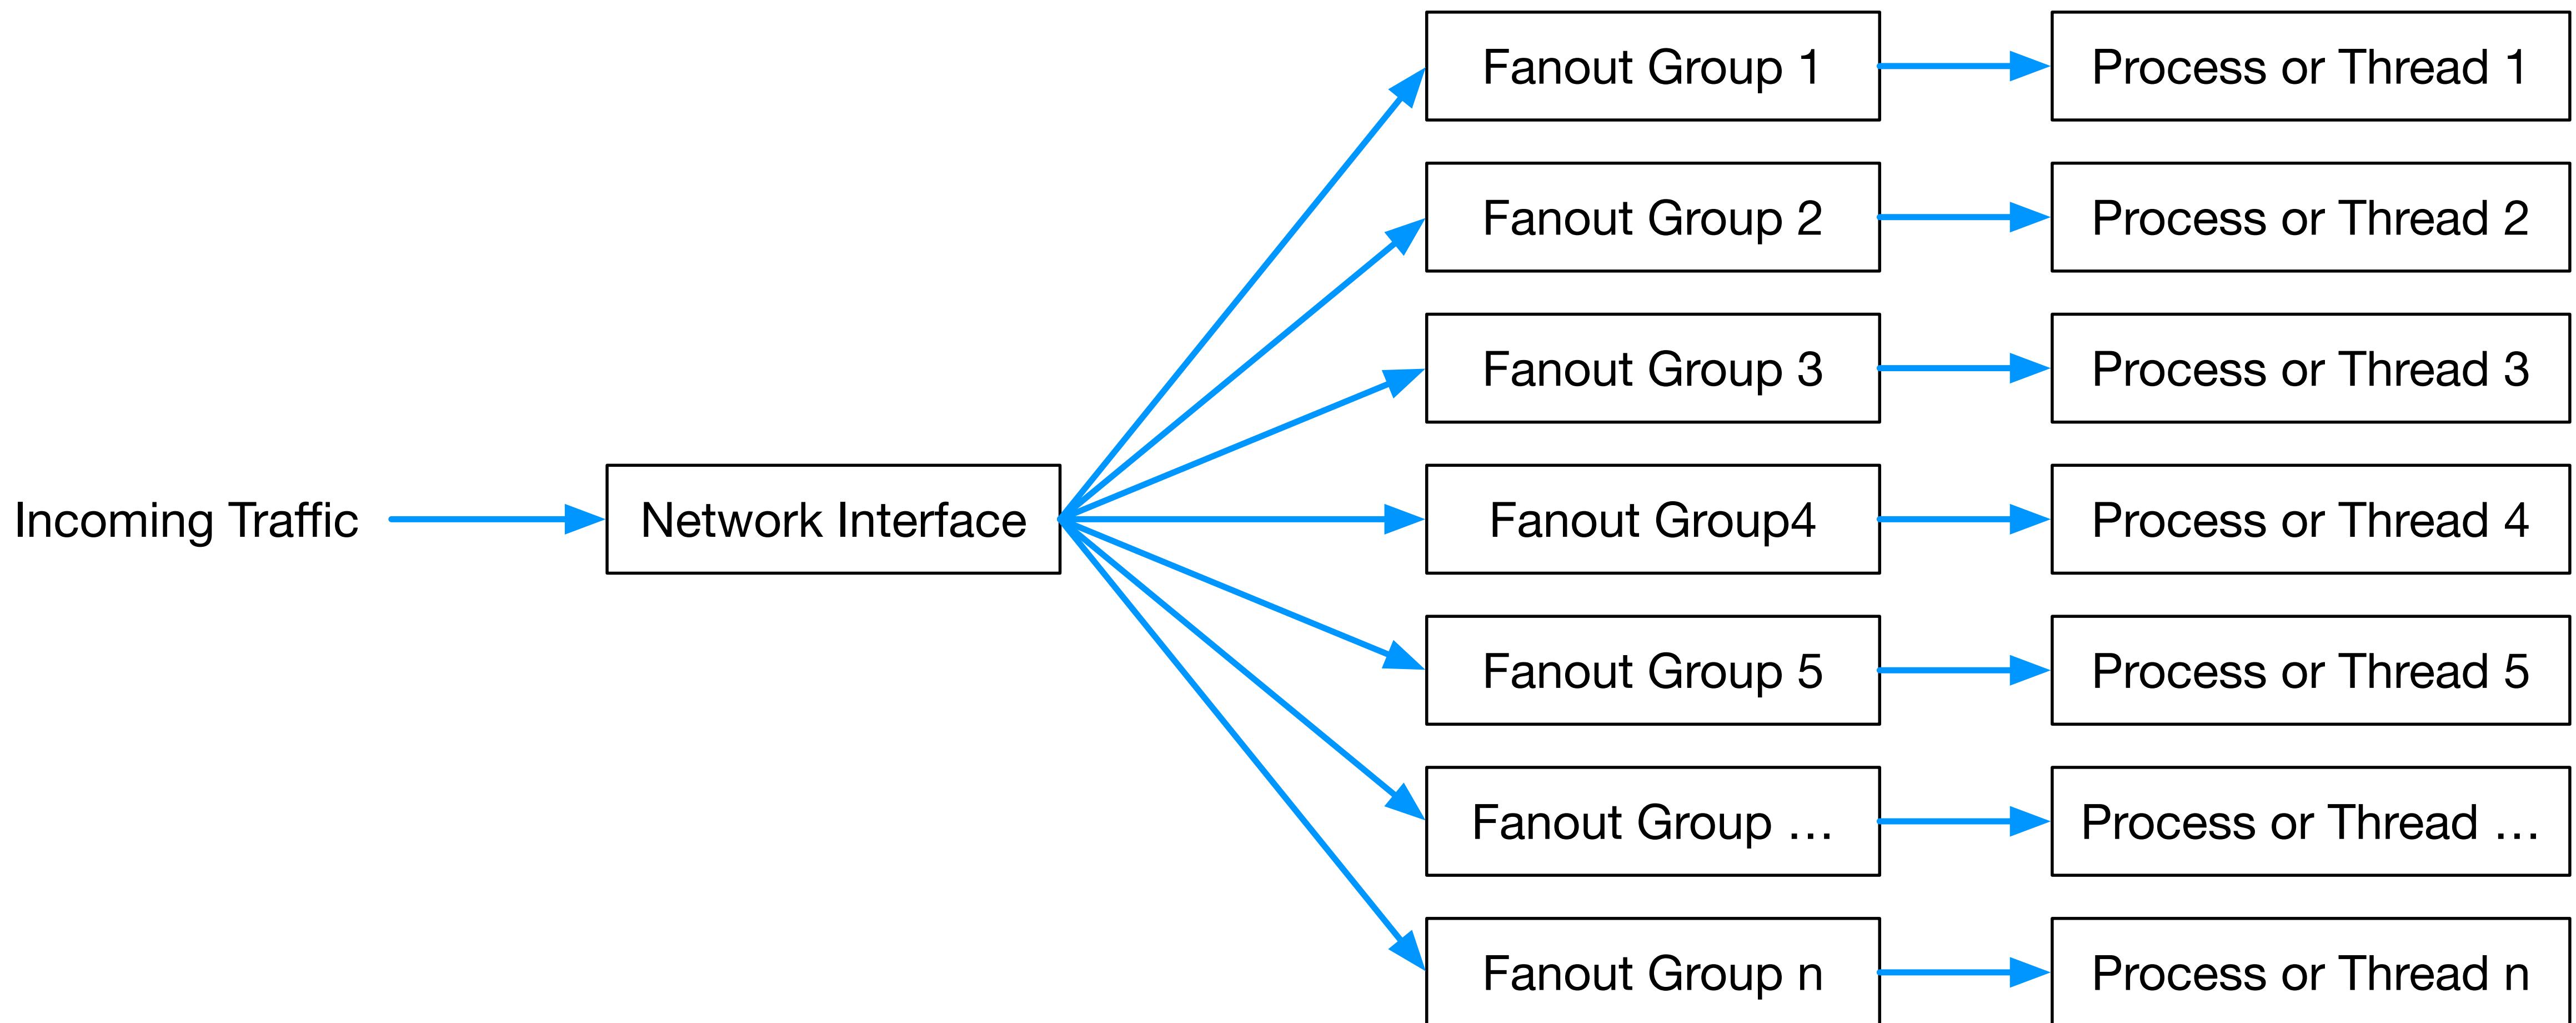




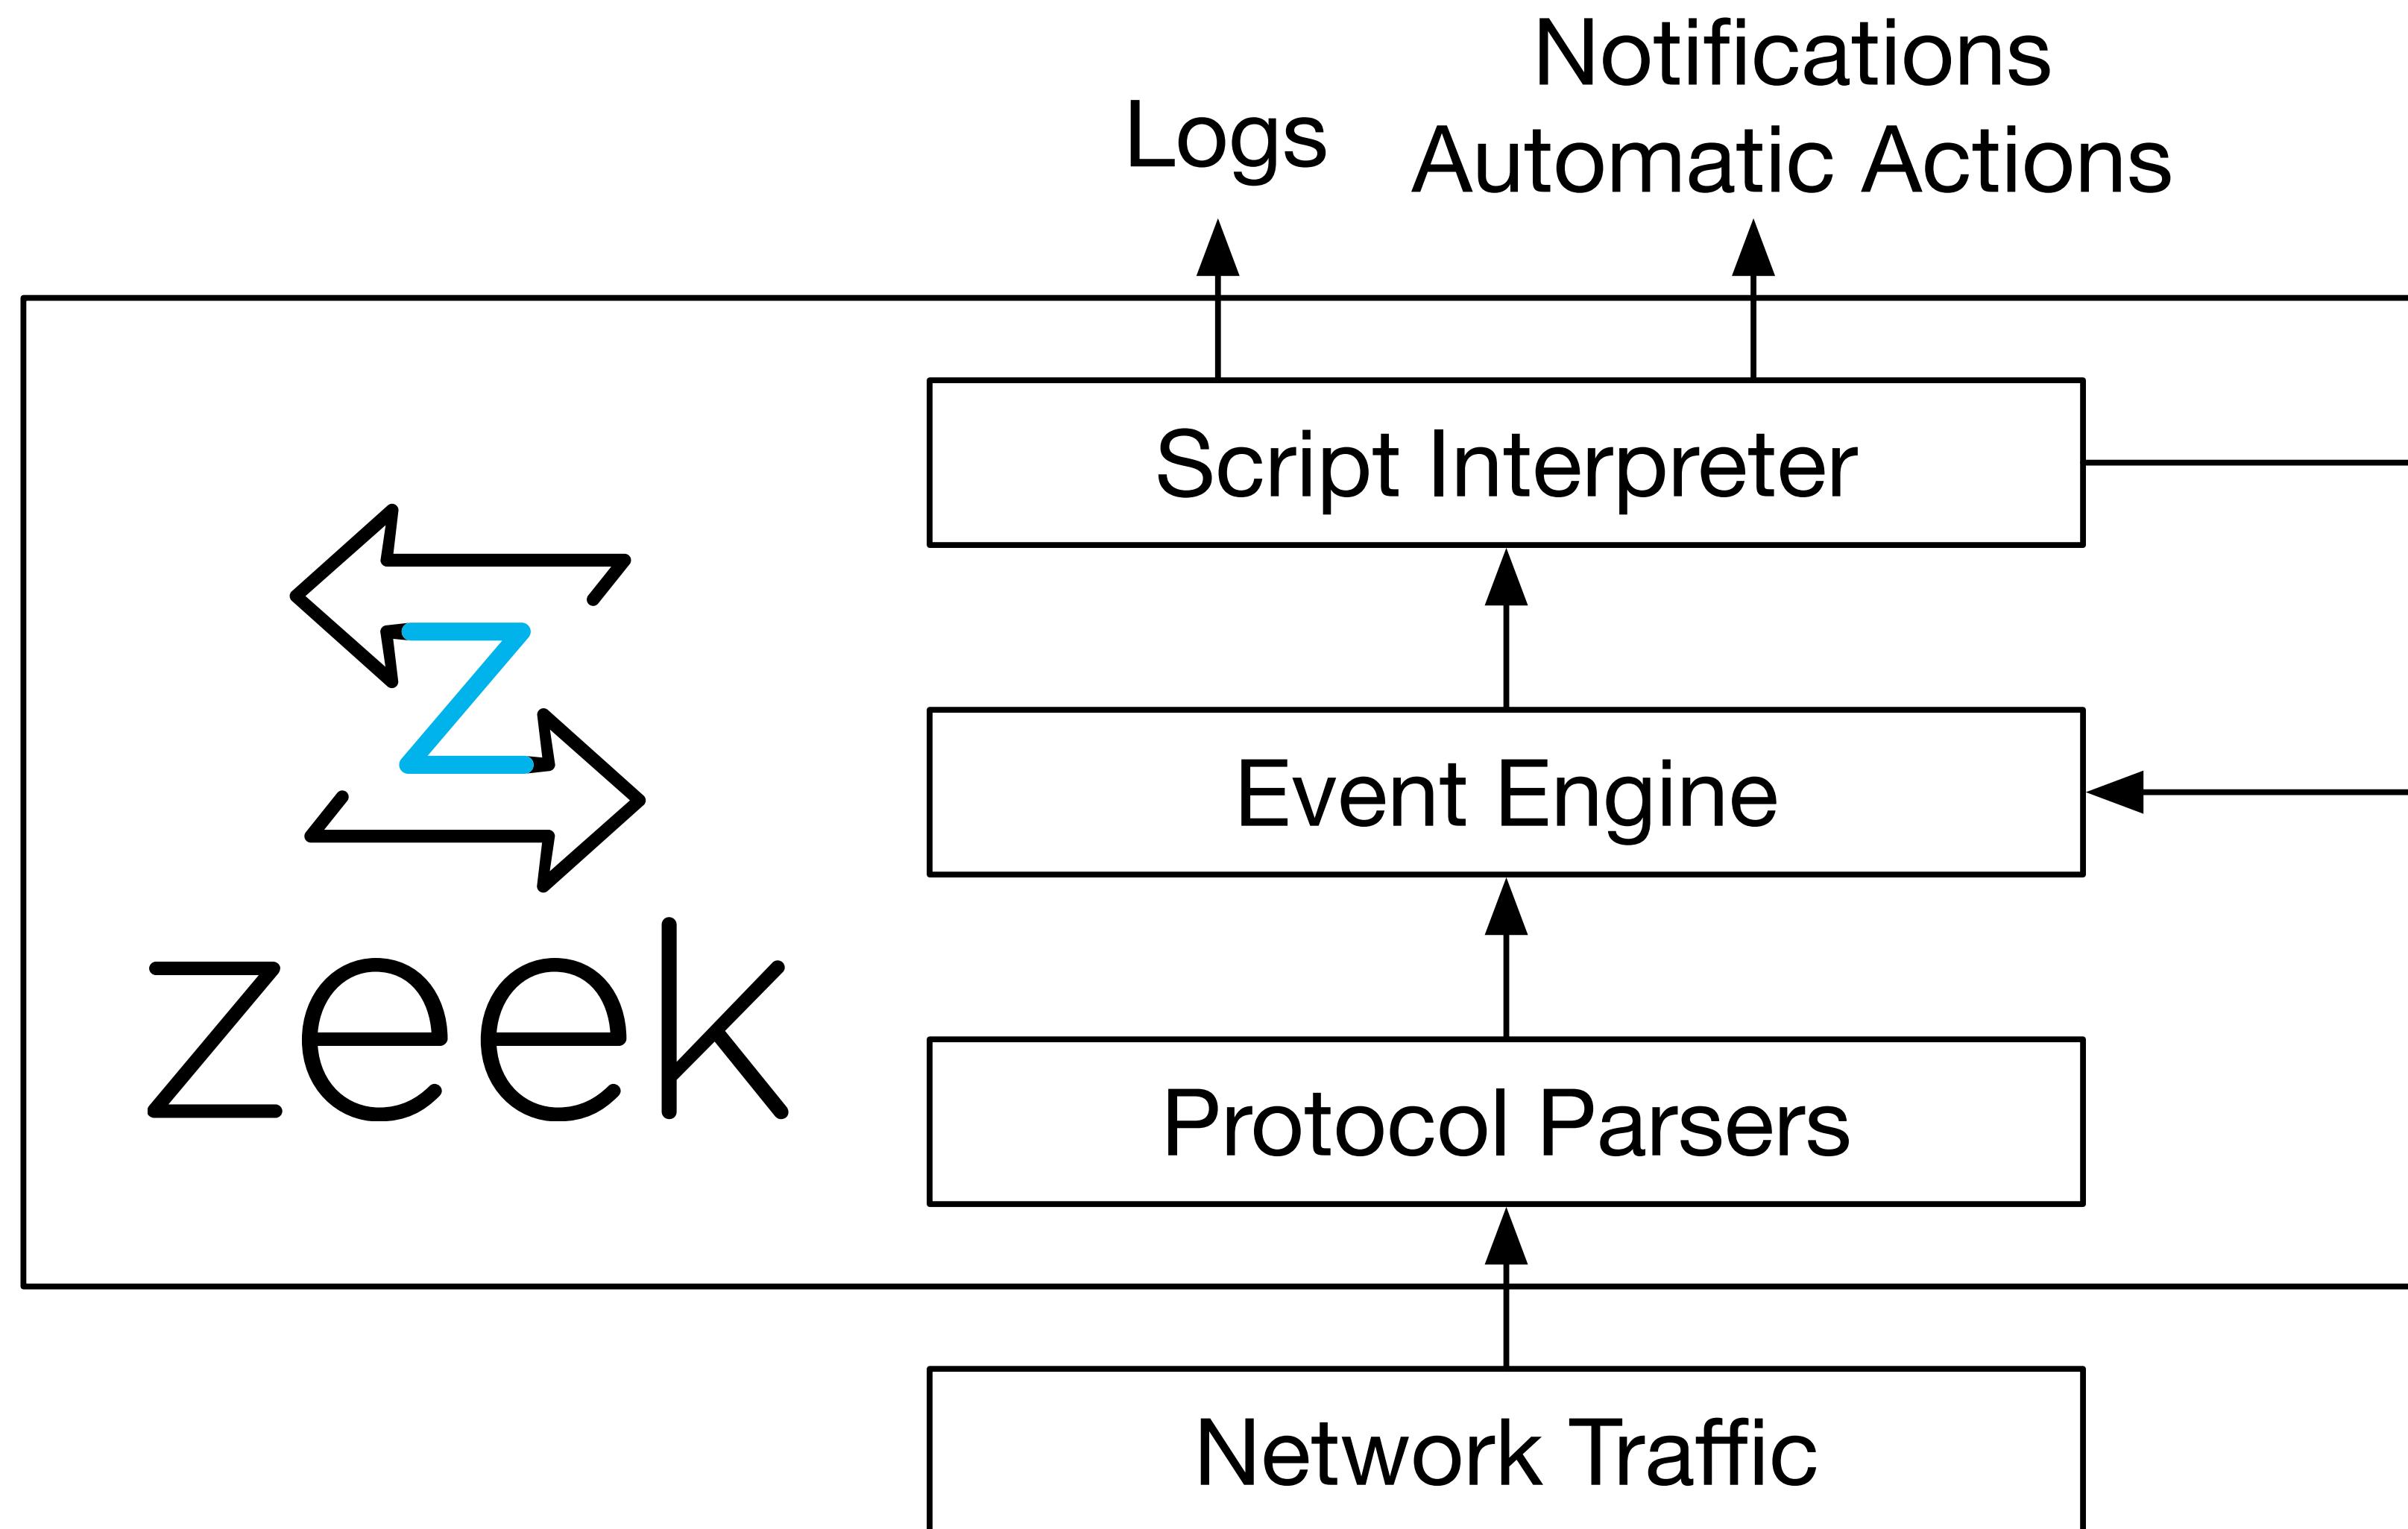




AF_PACKET Fanout



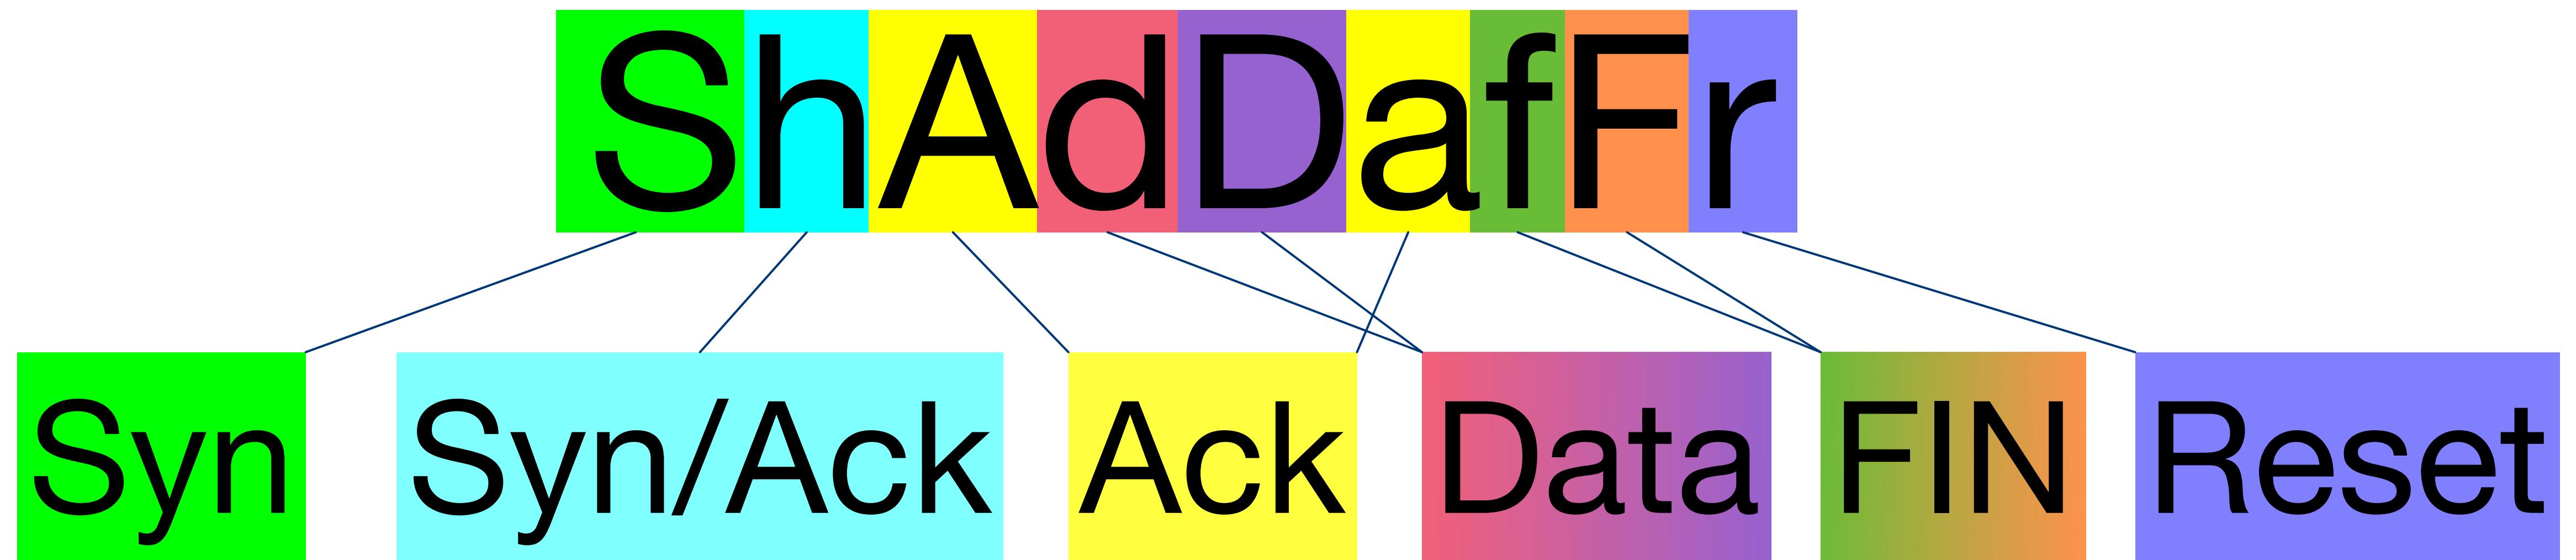
What is Zeek?



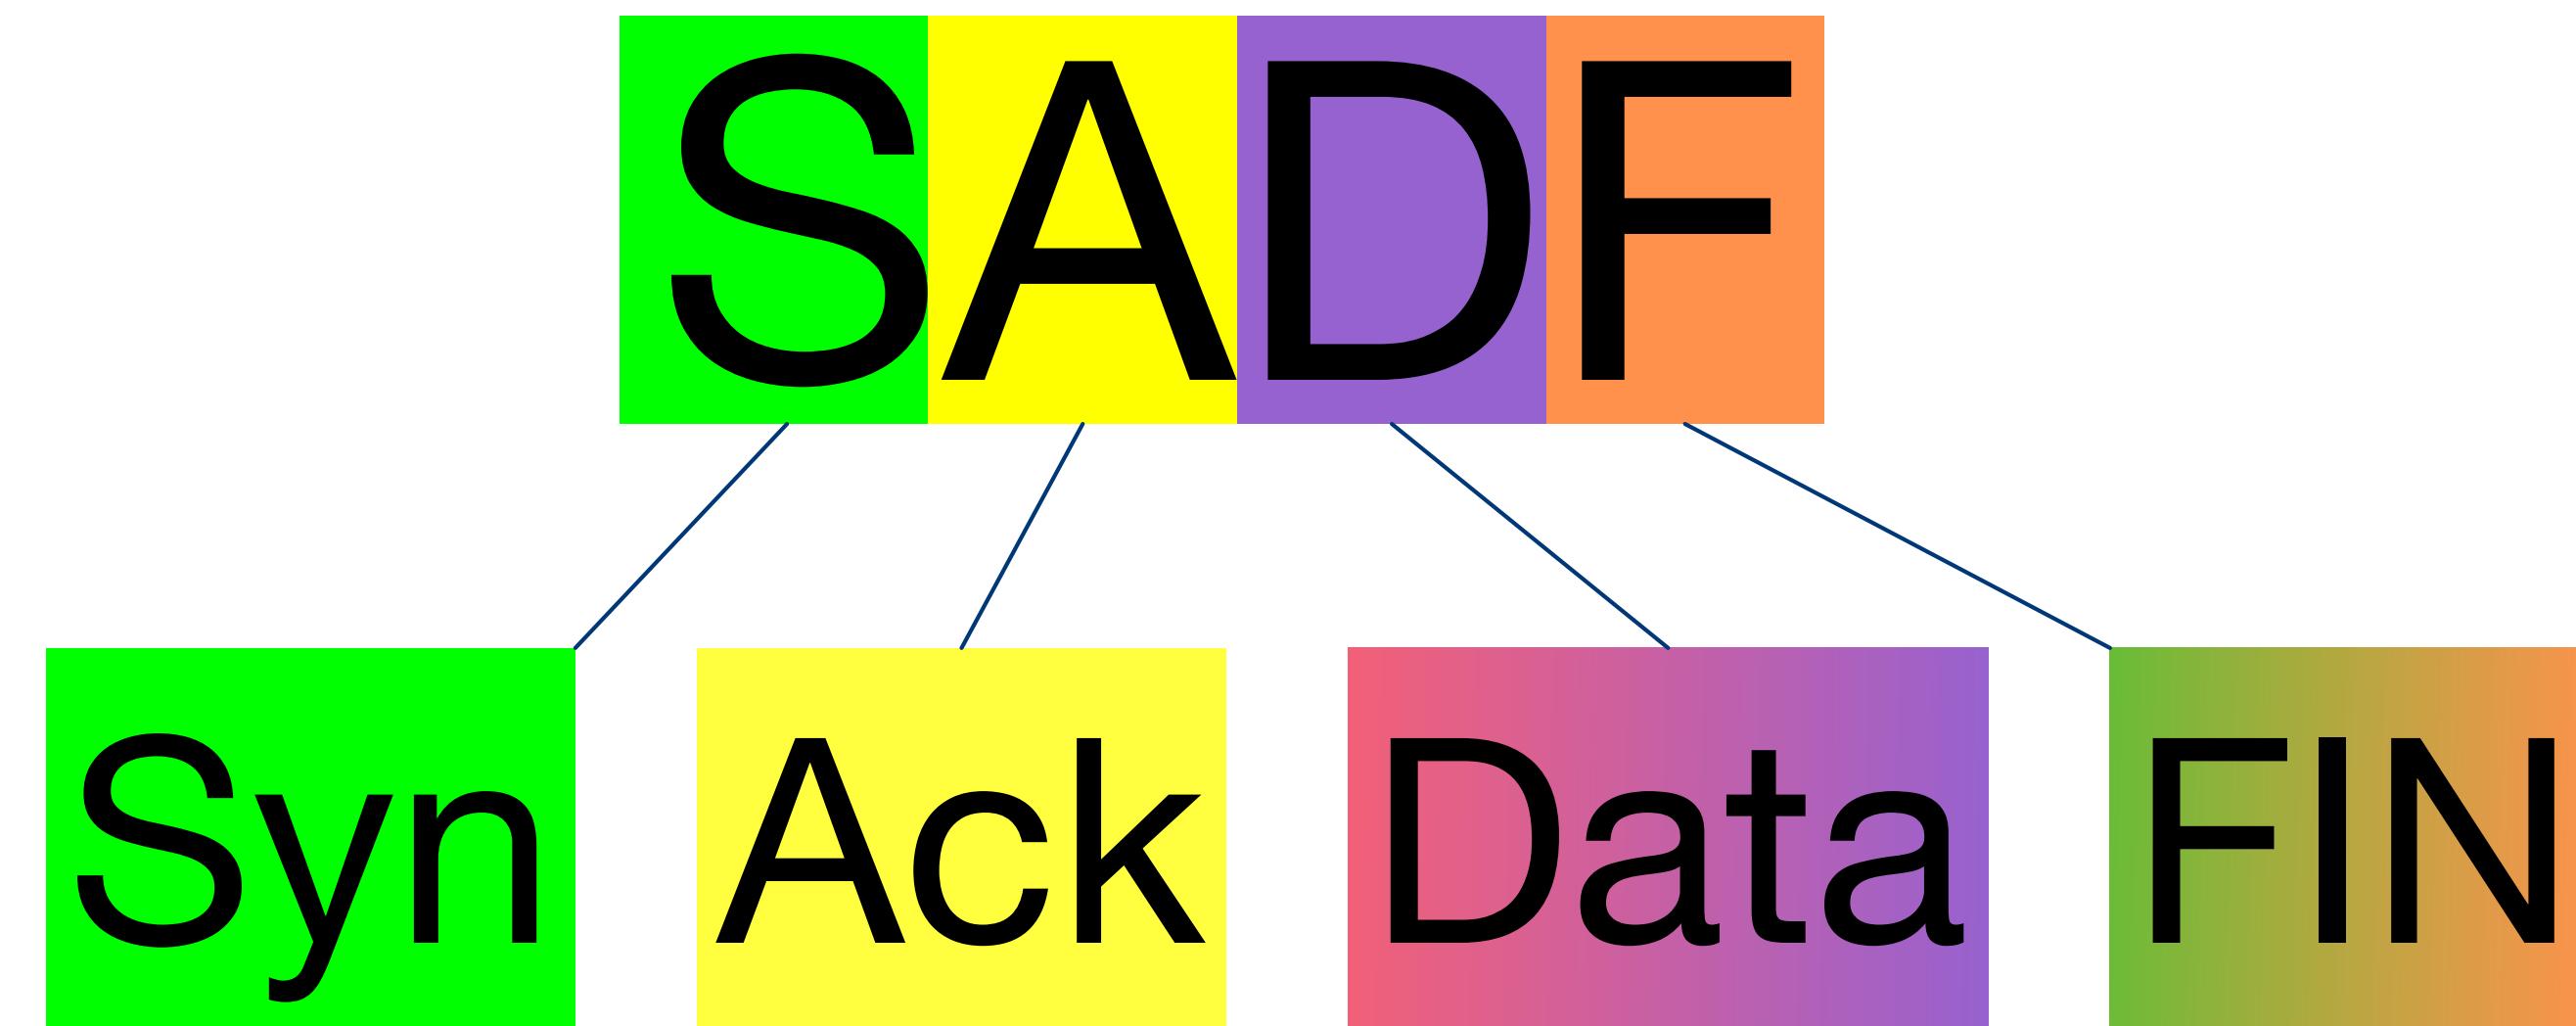
Check your data

- Goal: Be sure that your data is complete and internally consistent
- Often surprisingly tricky to do when you have a large amount of data
- First step: check that your TCP (and UDP) sessions look reasonable

Zeek Connection History



Zeek Connection History



Check that your data is complete

Ethernet

Vlan 10

Data

Check that your data is complete

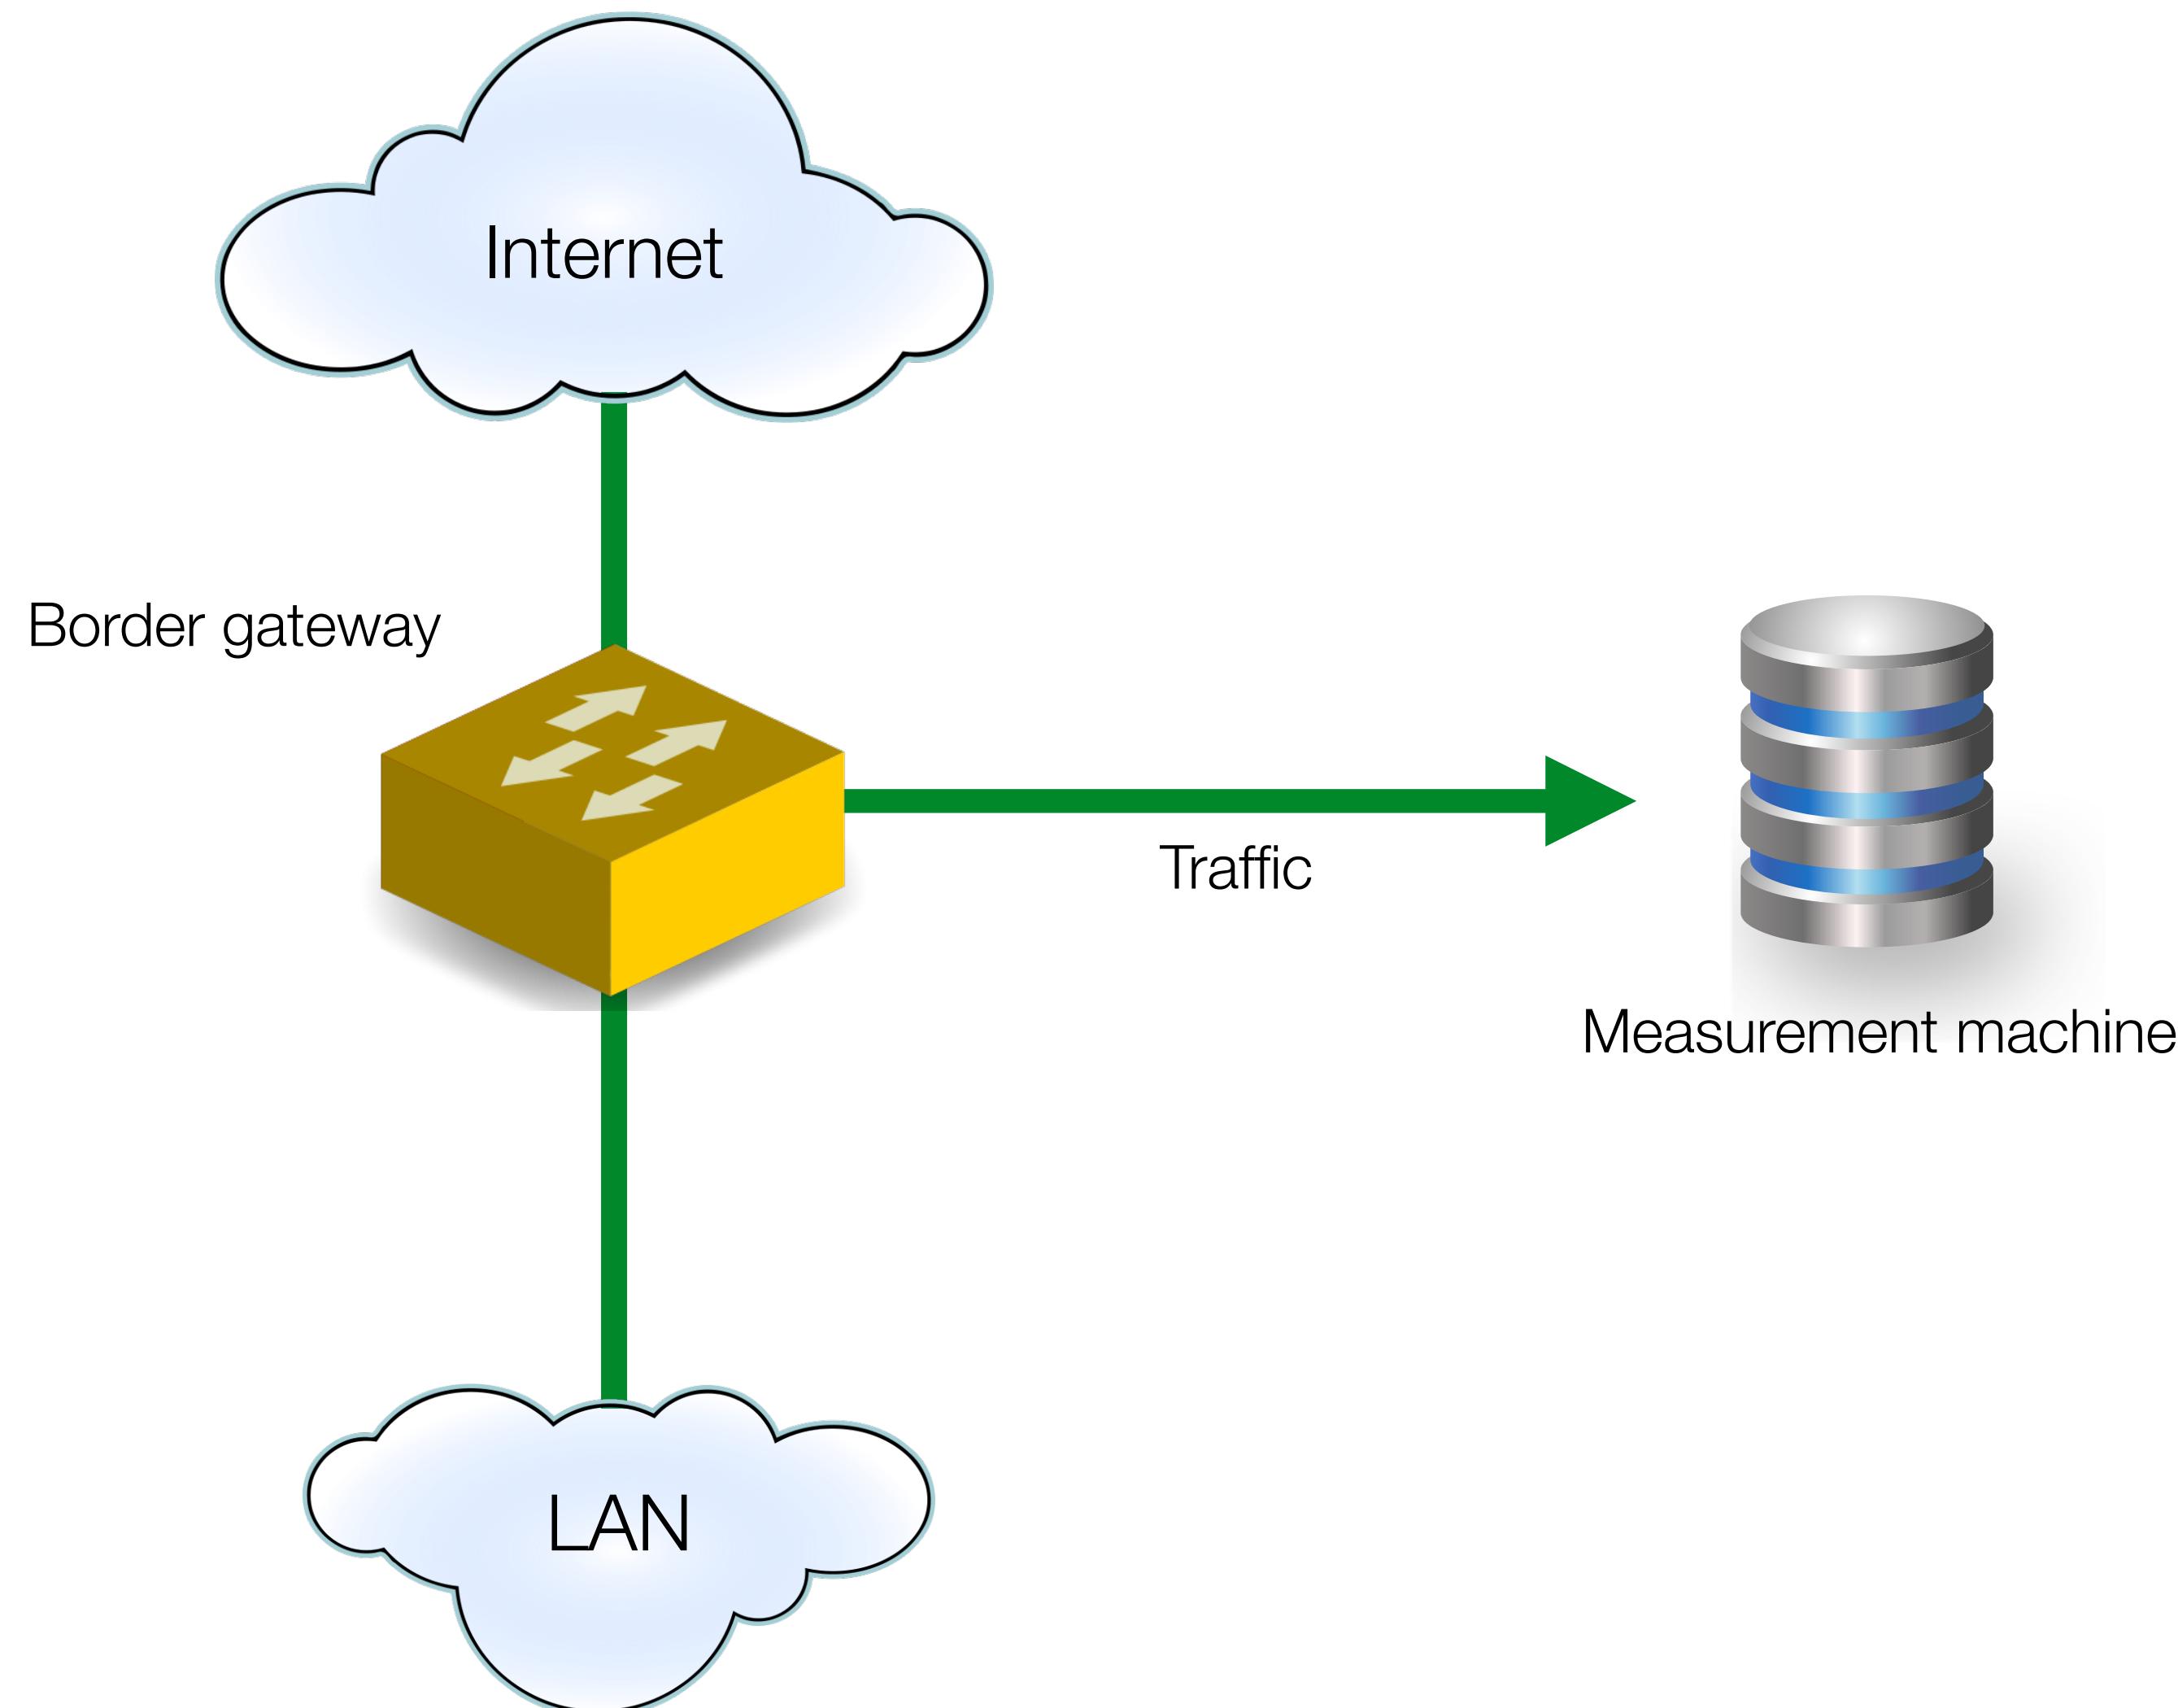
Ethernet

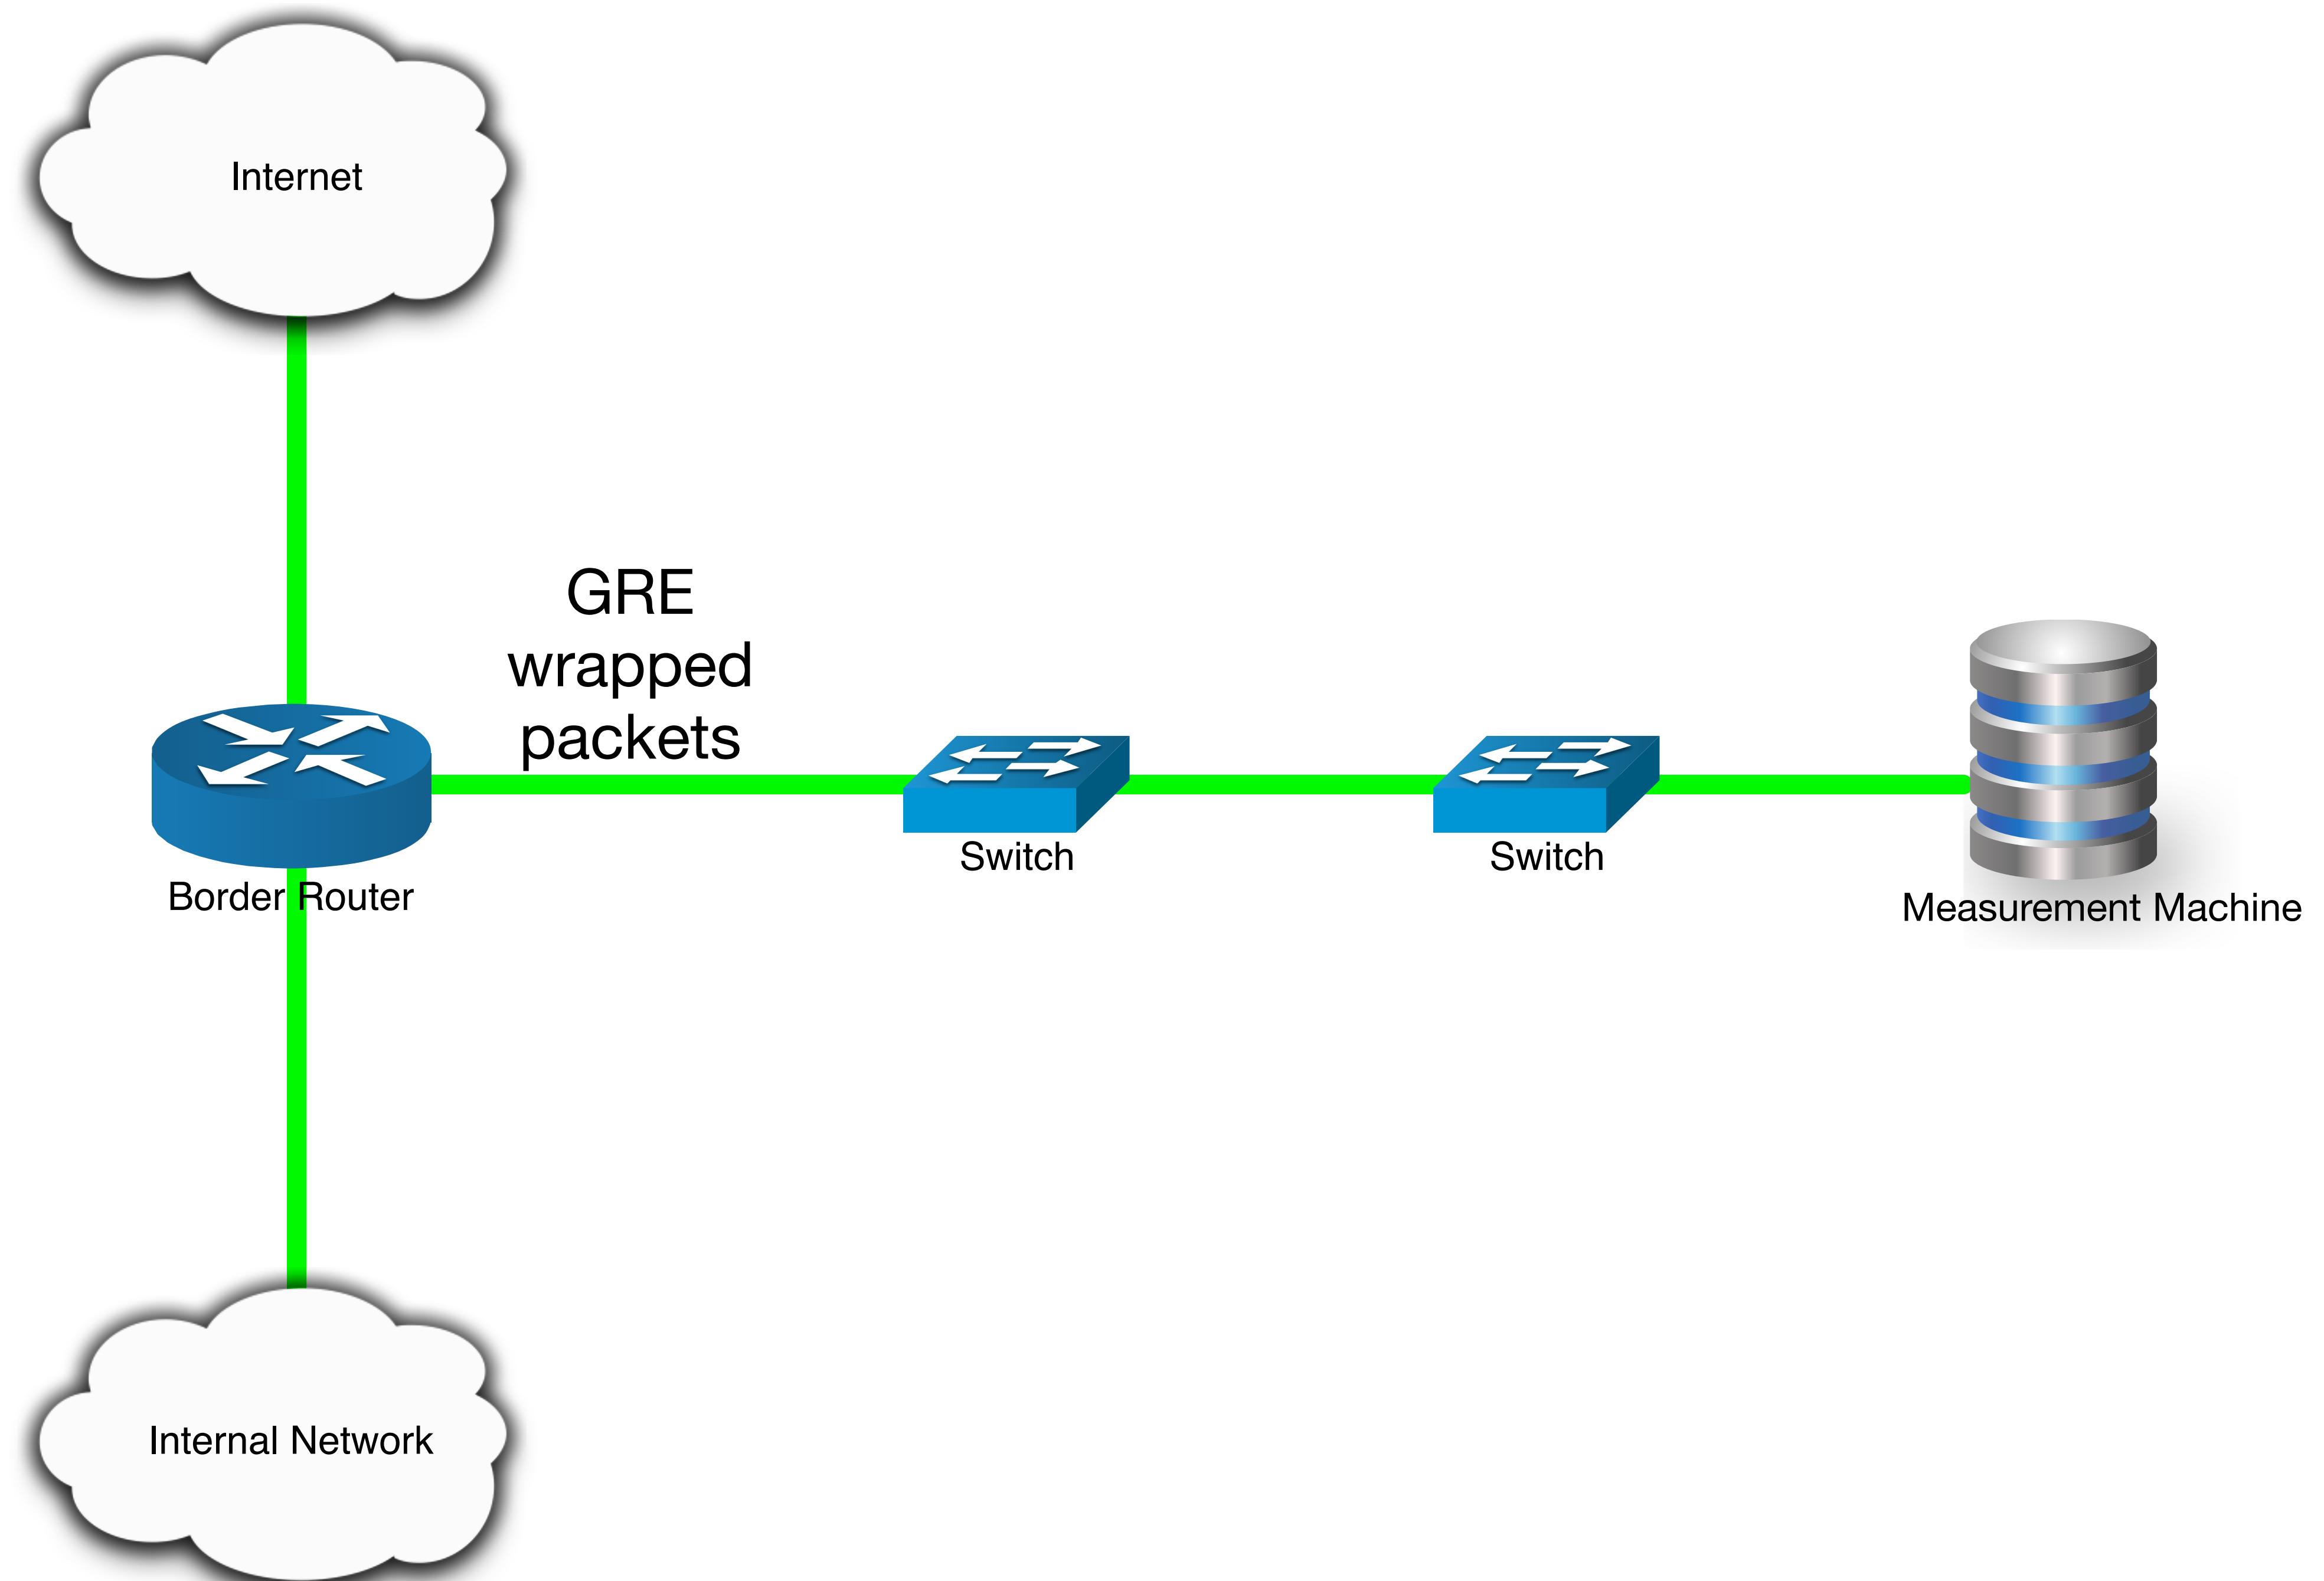
Vlan 10

Vlan 20

Data

Typical Network Monitoring Setup

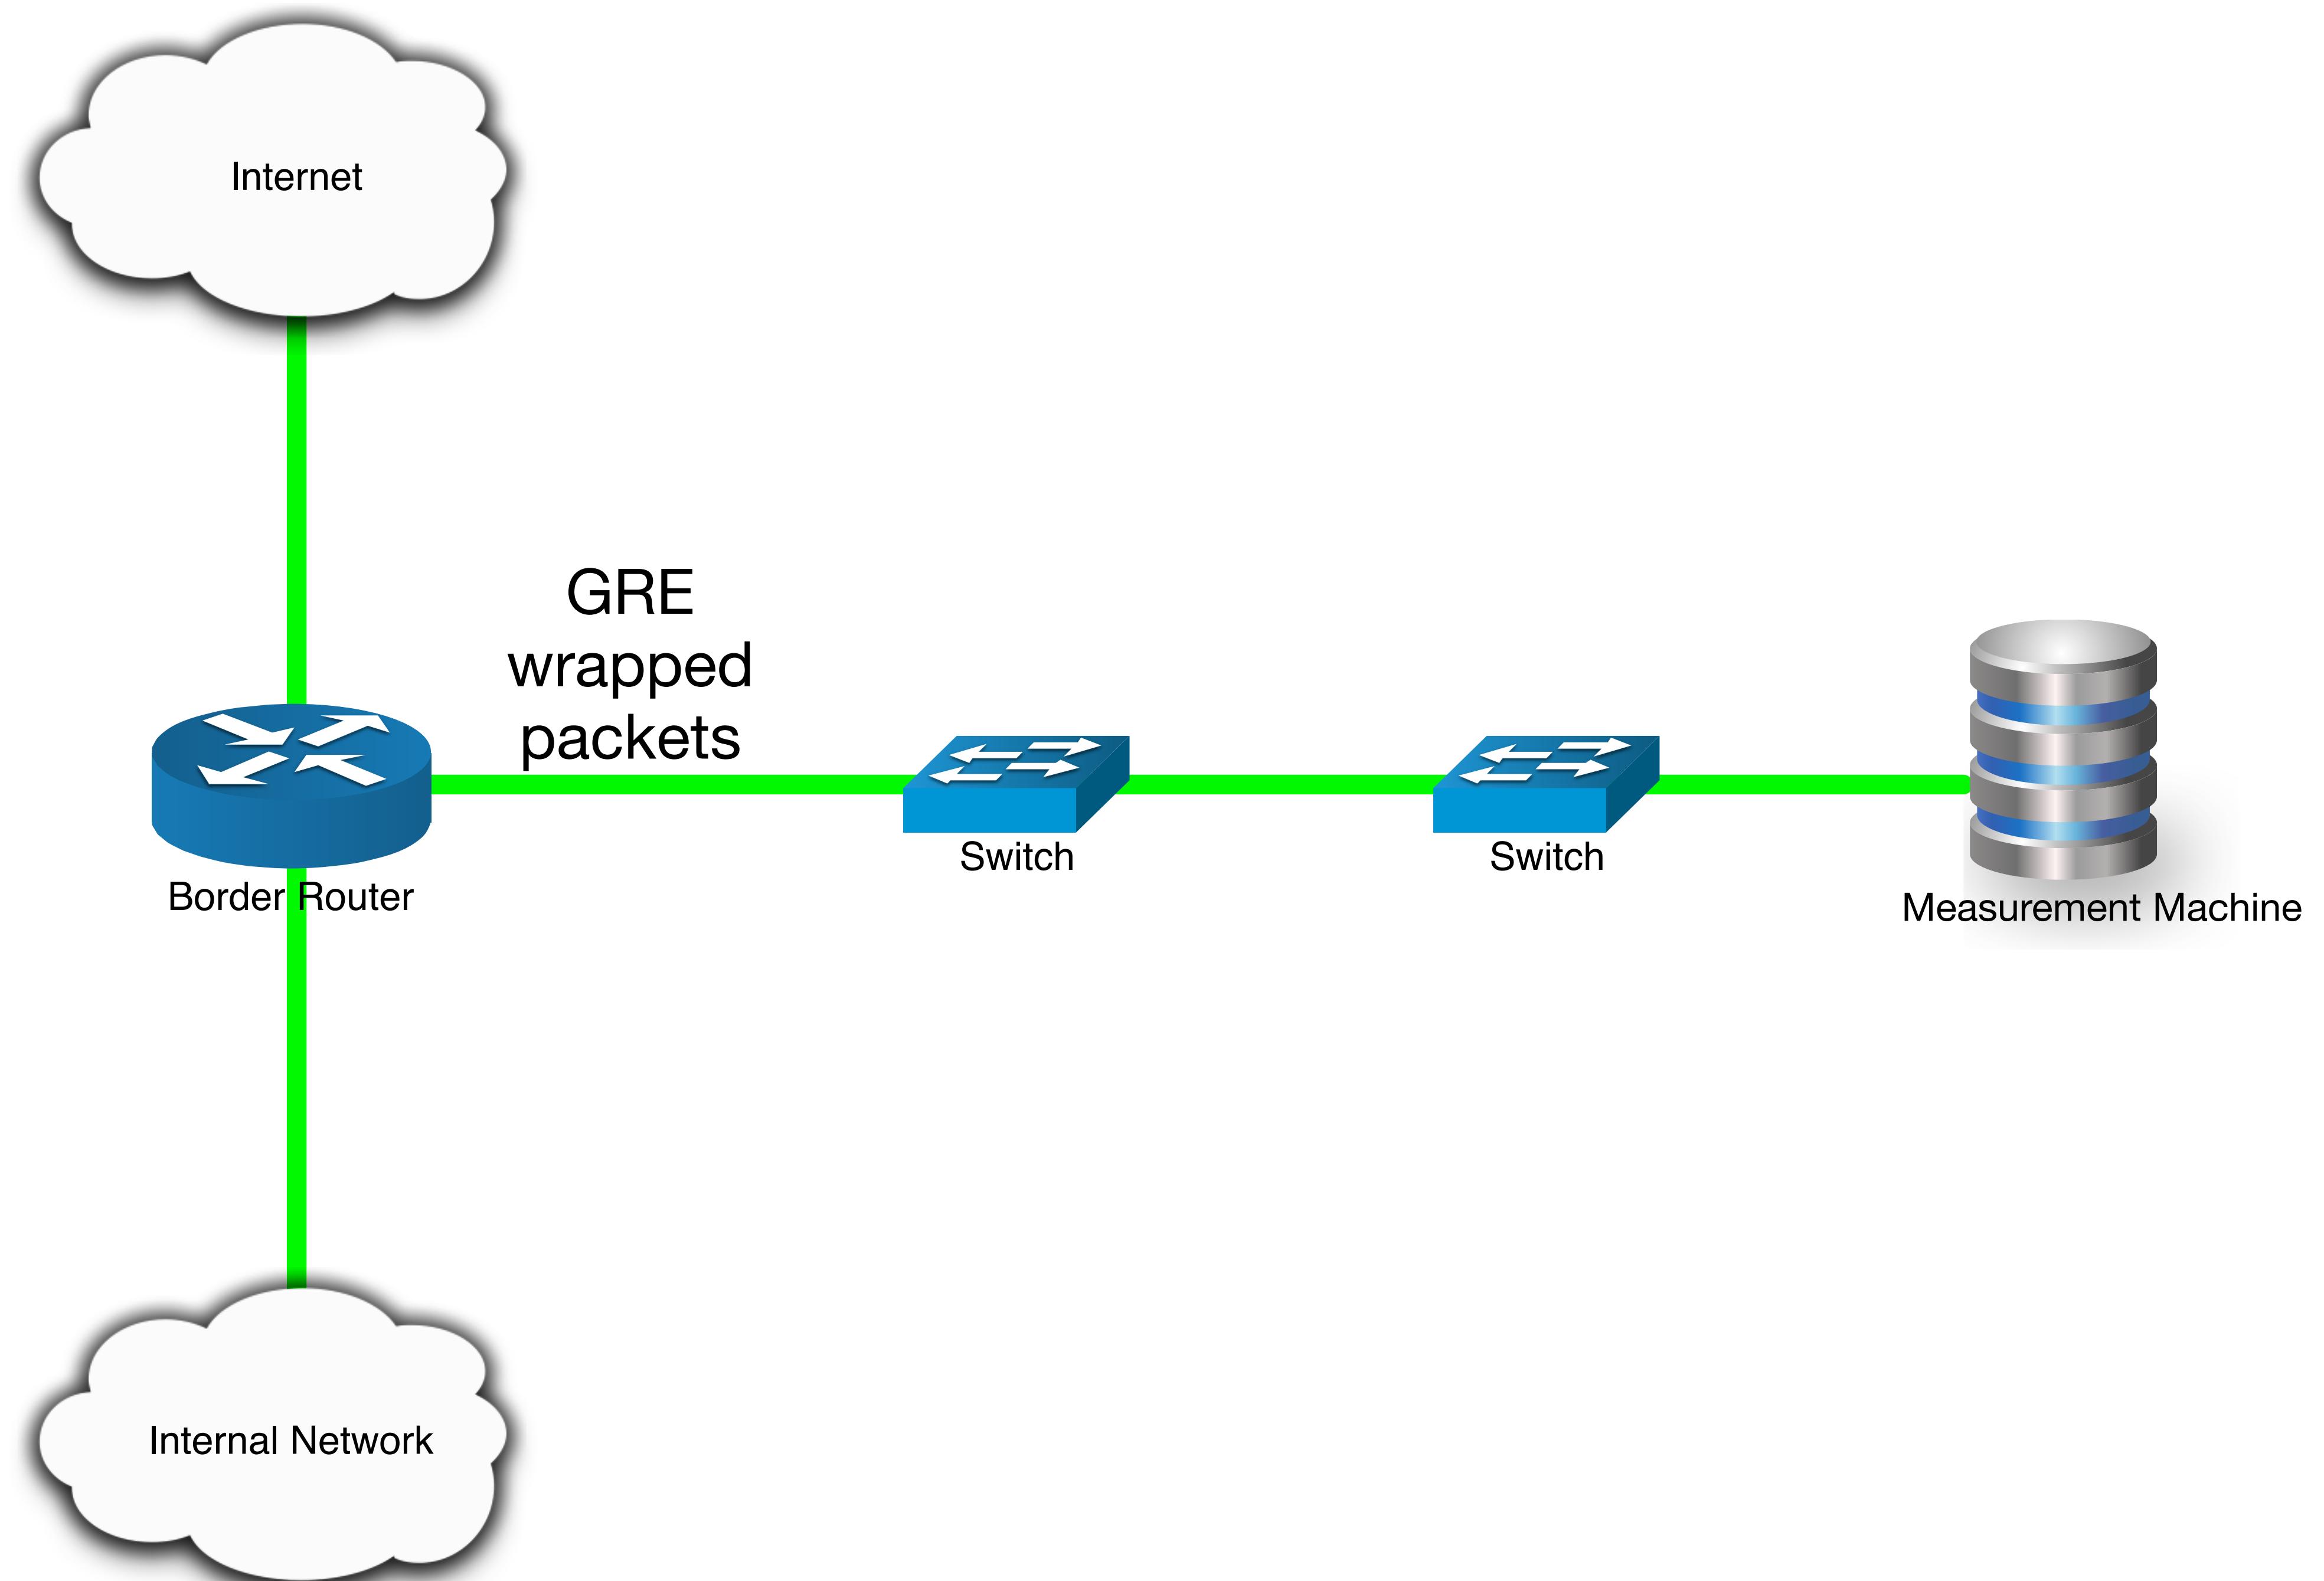




Capture loss

- Zeek log that uses TCP Gap/ACK analysis to determine actual amount of data lost

ts_delta	Gaps	Acks	Perecent_lost
900	93372	4593277	2.04



check your data - and your pipeline

How do you get contributors?



Getting contributors

- Asking gets you surprisingly far
- Operations often are interesting in helping you
- It pays off to have good long-term relationships
- People that do this operationally already have monitoring, hardware, etc
So you get the data without having to spend your research money on it
- But - you don't control the data pipeline - so you can't guarantee its accuracy or bias
- Software updates are hard
- Think about schemas early
Make them extensible
You probably still will end up with a mess :)

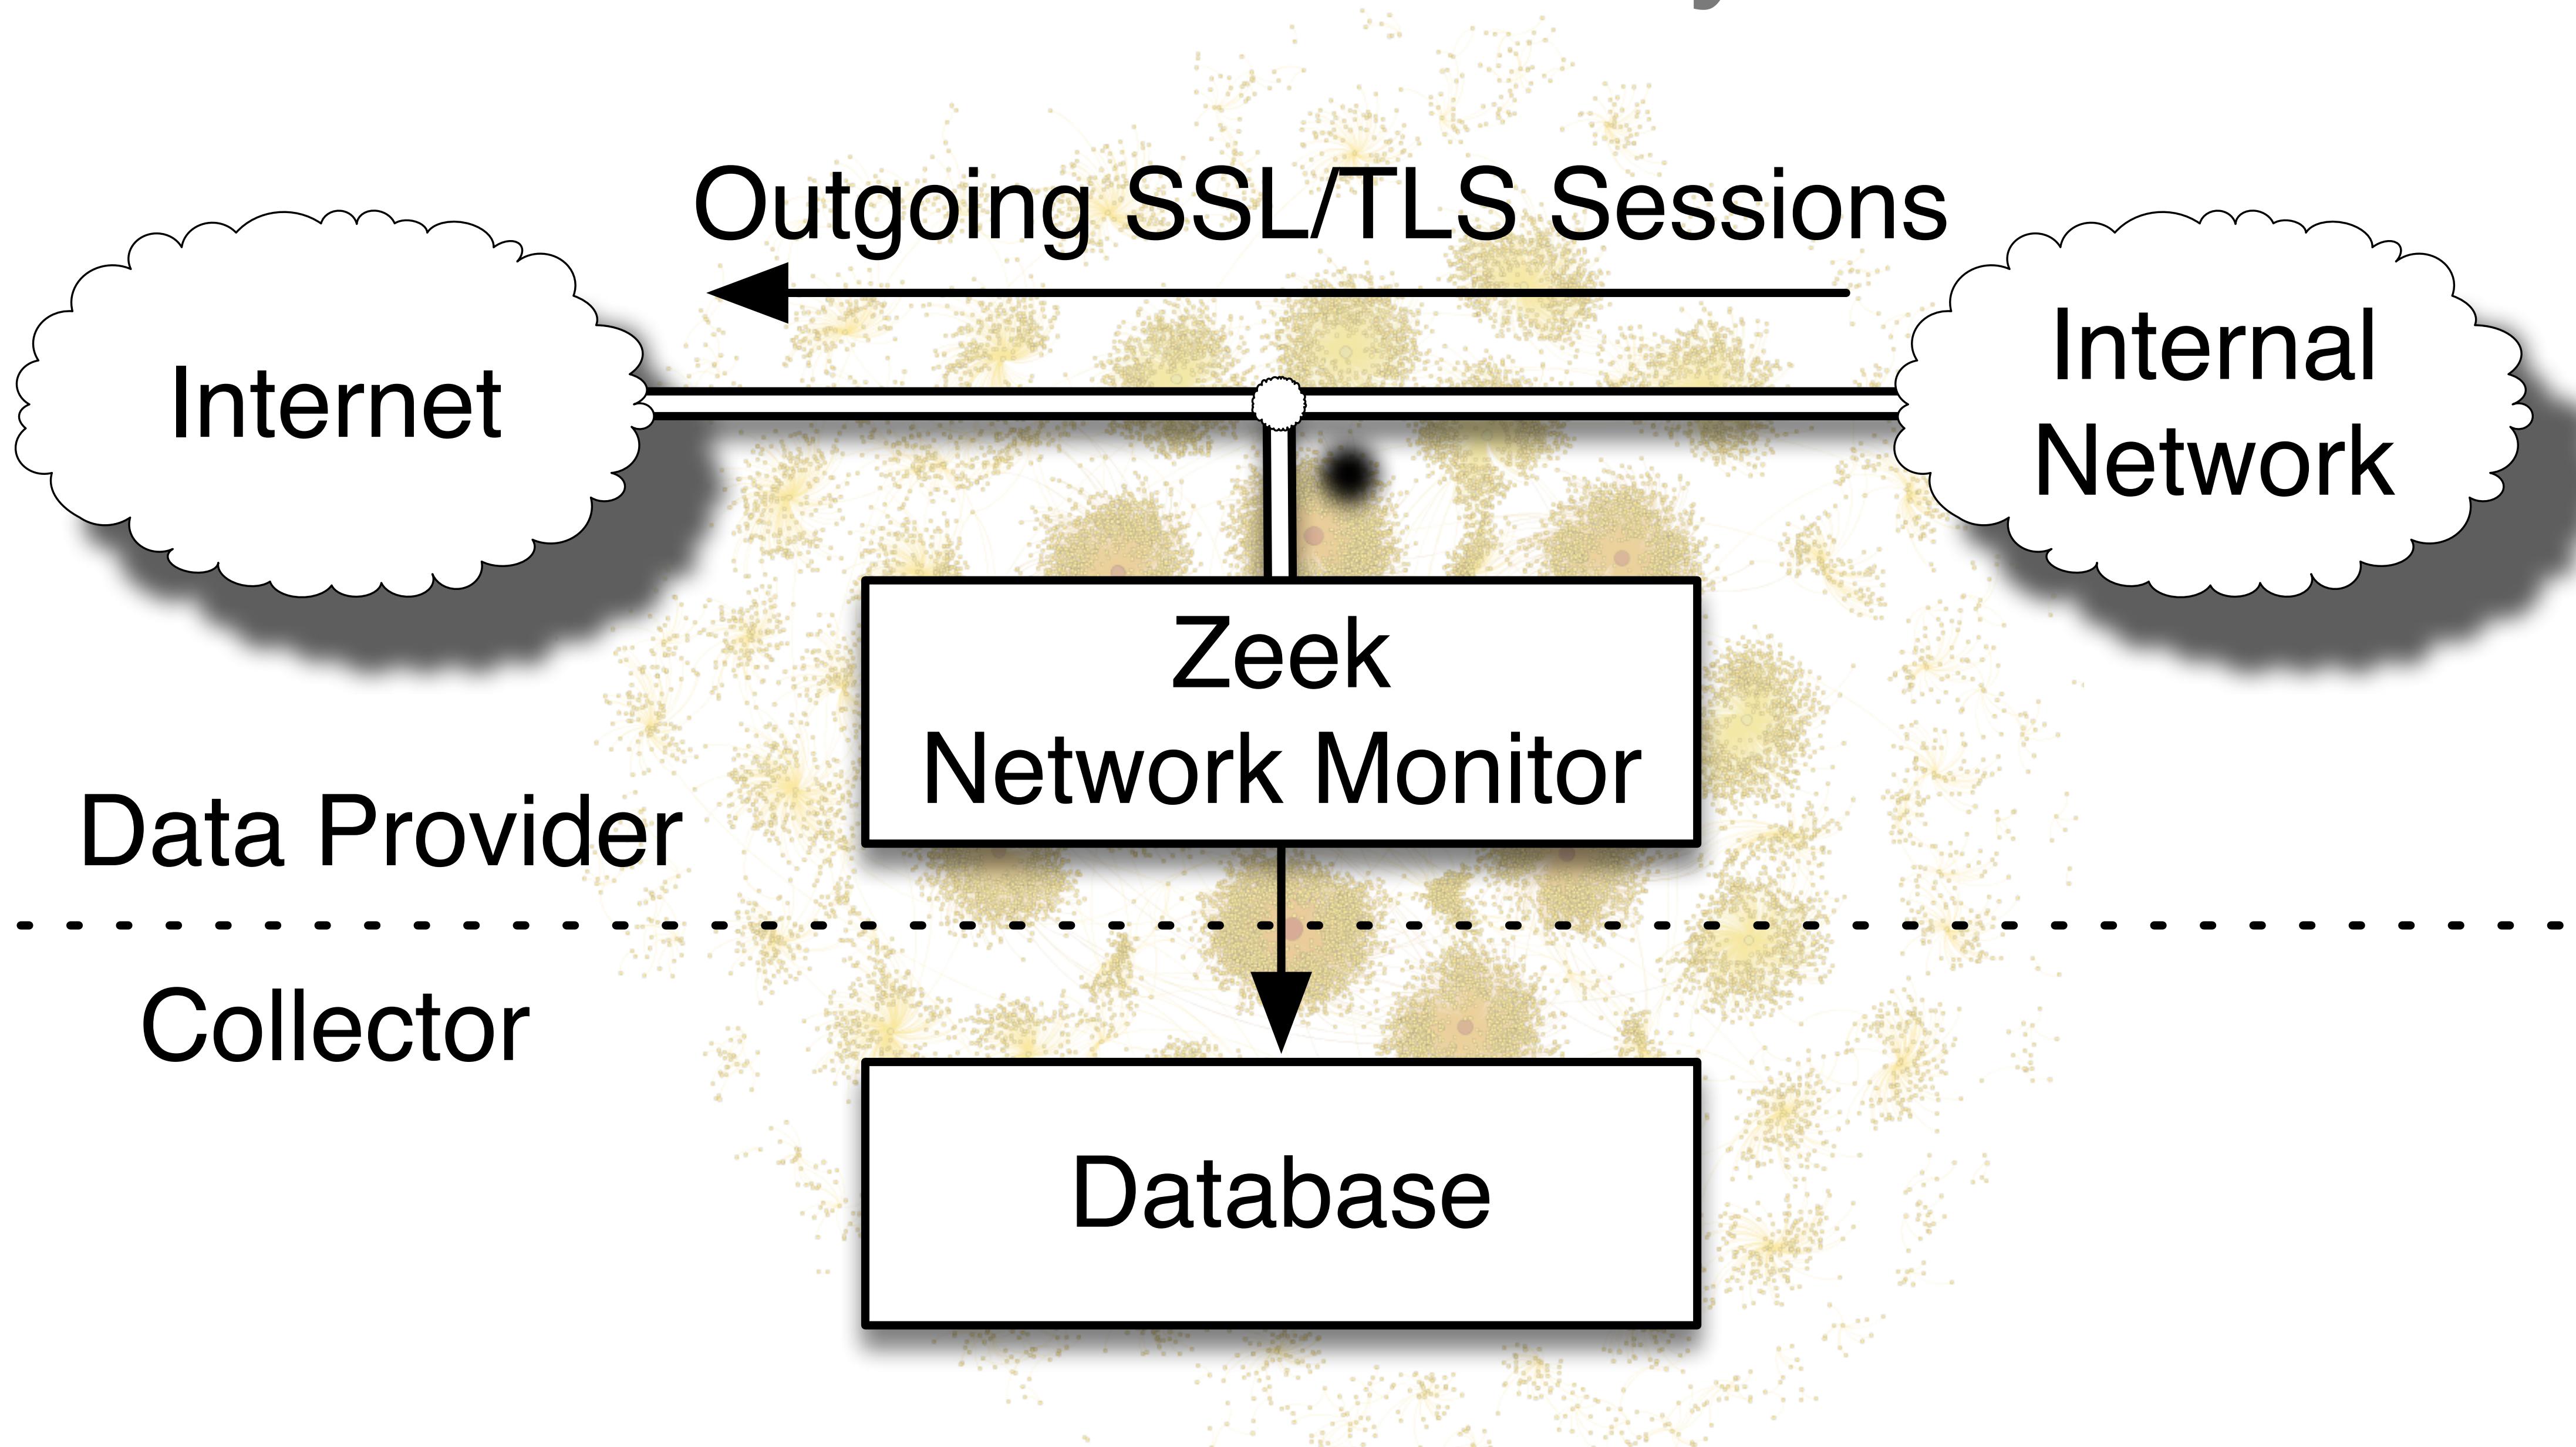
Lessons so far

- Setting passive measurements up correctly takes time
- Think about ethics early
- Validate your measurements - and continue re-checking them.
Things can (and will) break.
- Debugging these issues can take a lot of time
It is not always possible to fix things easily, even when you identified the issue
- External contributors are great - but you don't control the pipeline

Part 2

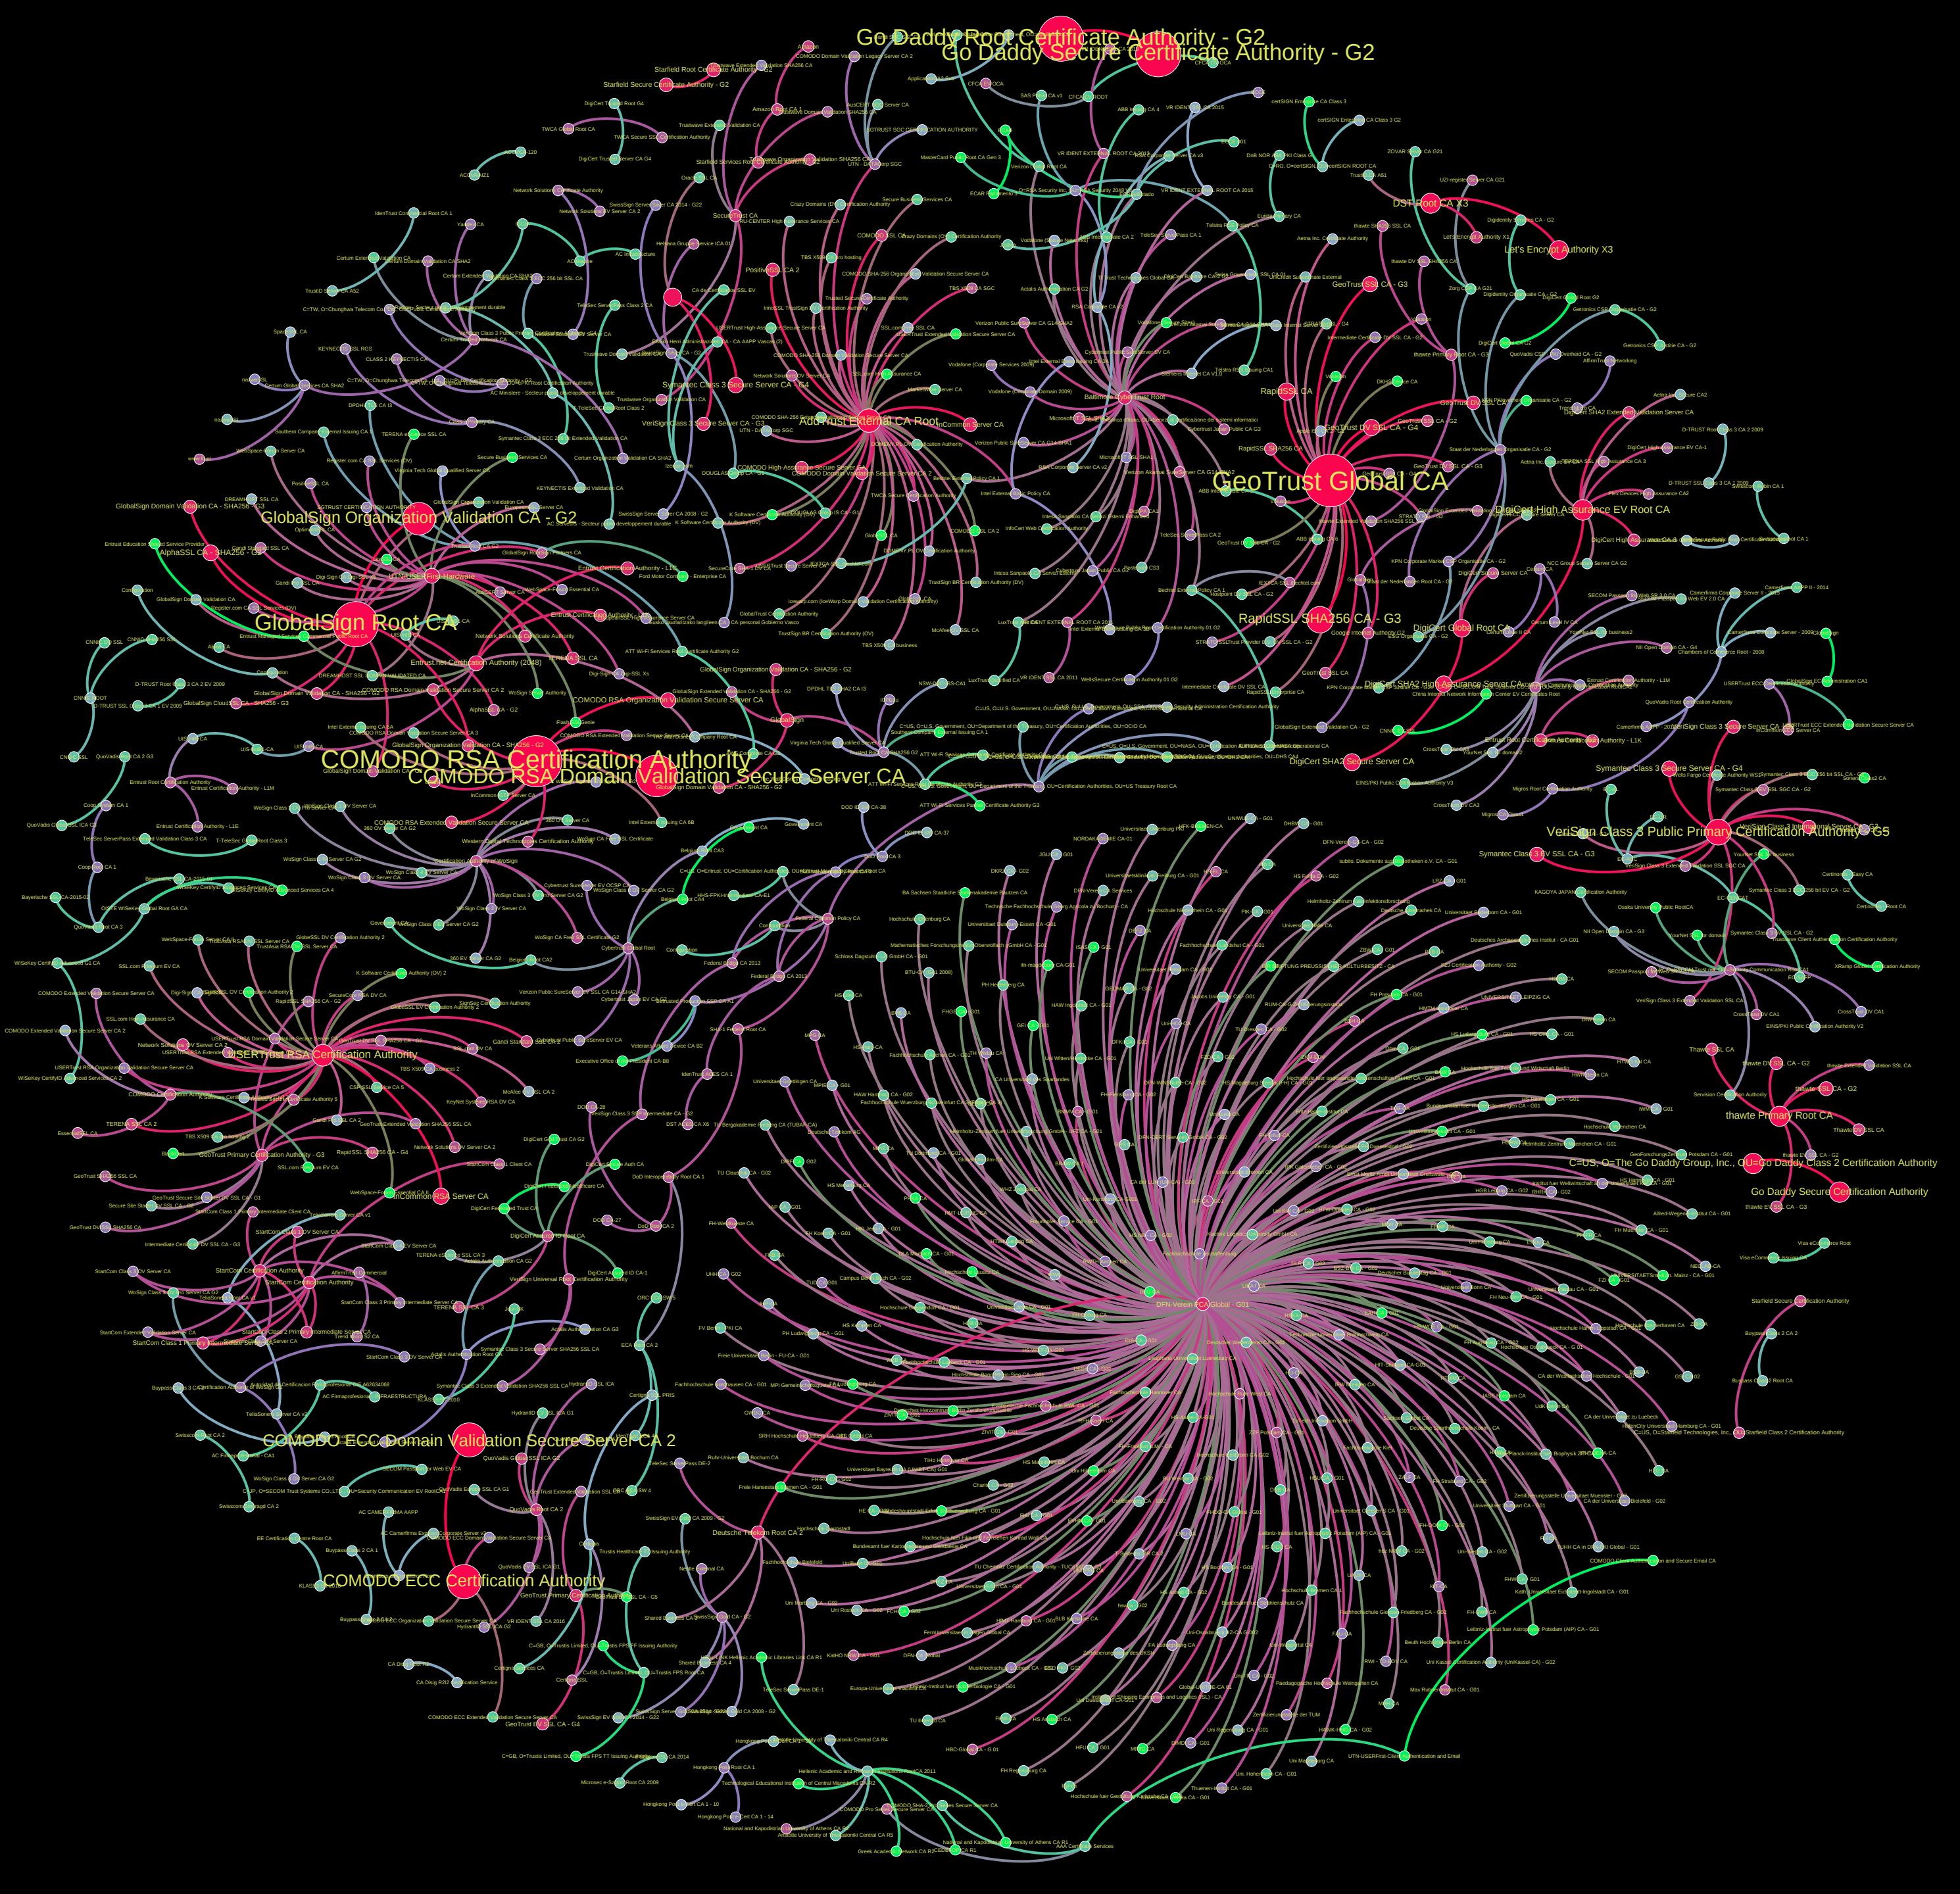
TLS and Protocol Weirdnesses

ICSI Notary

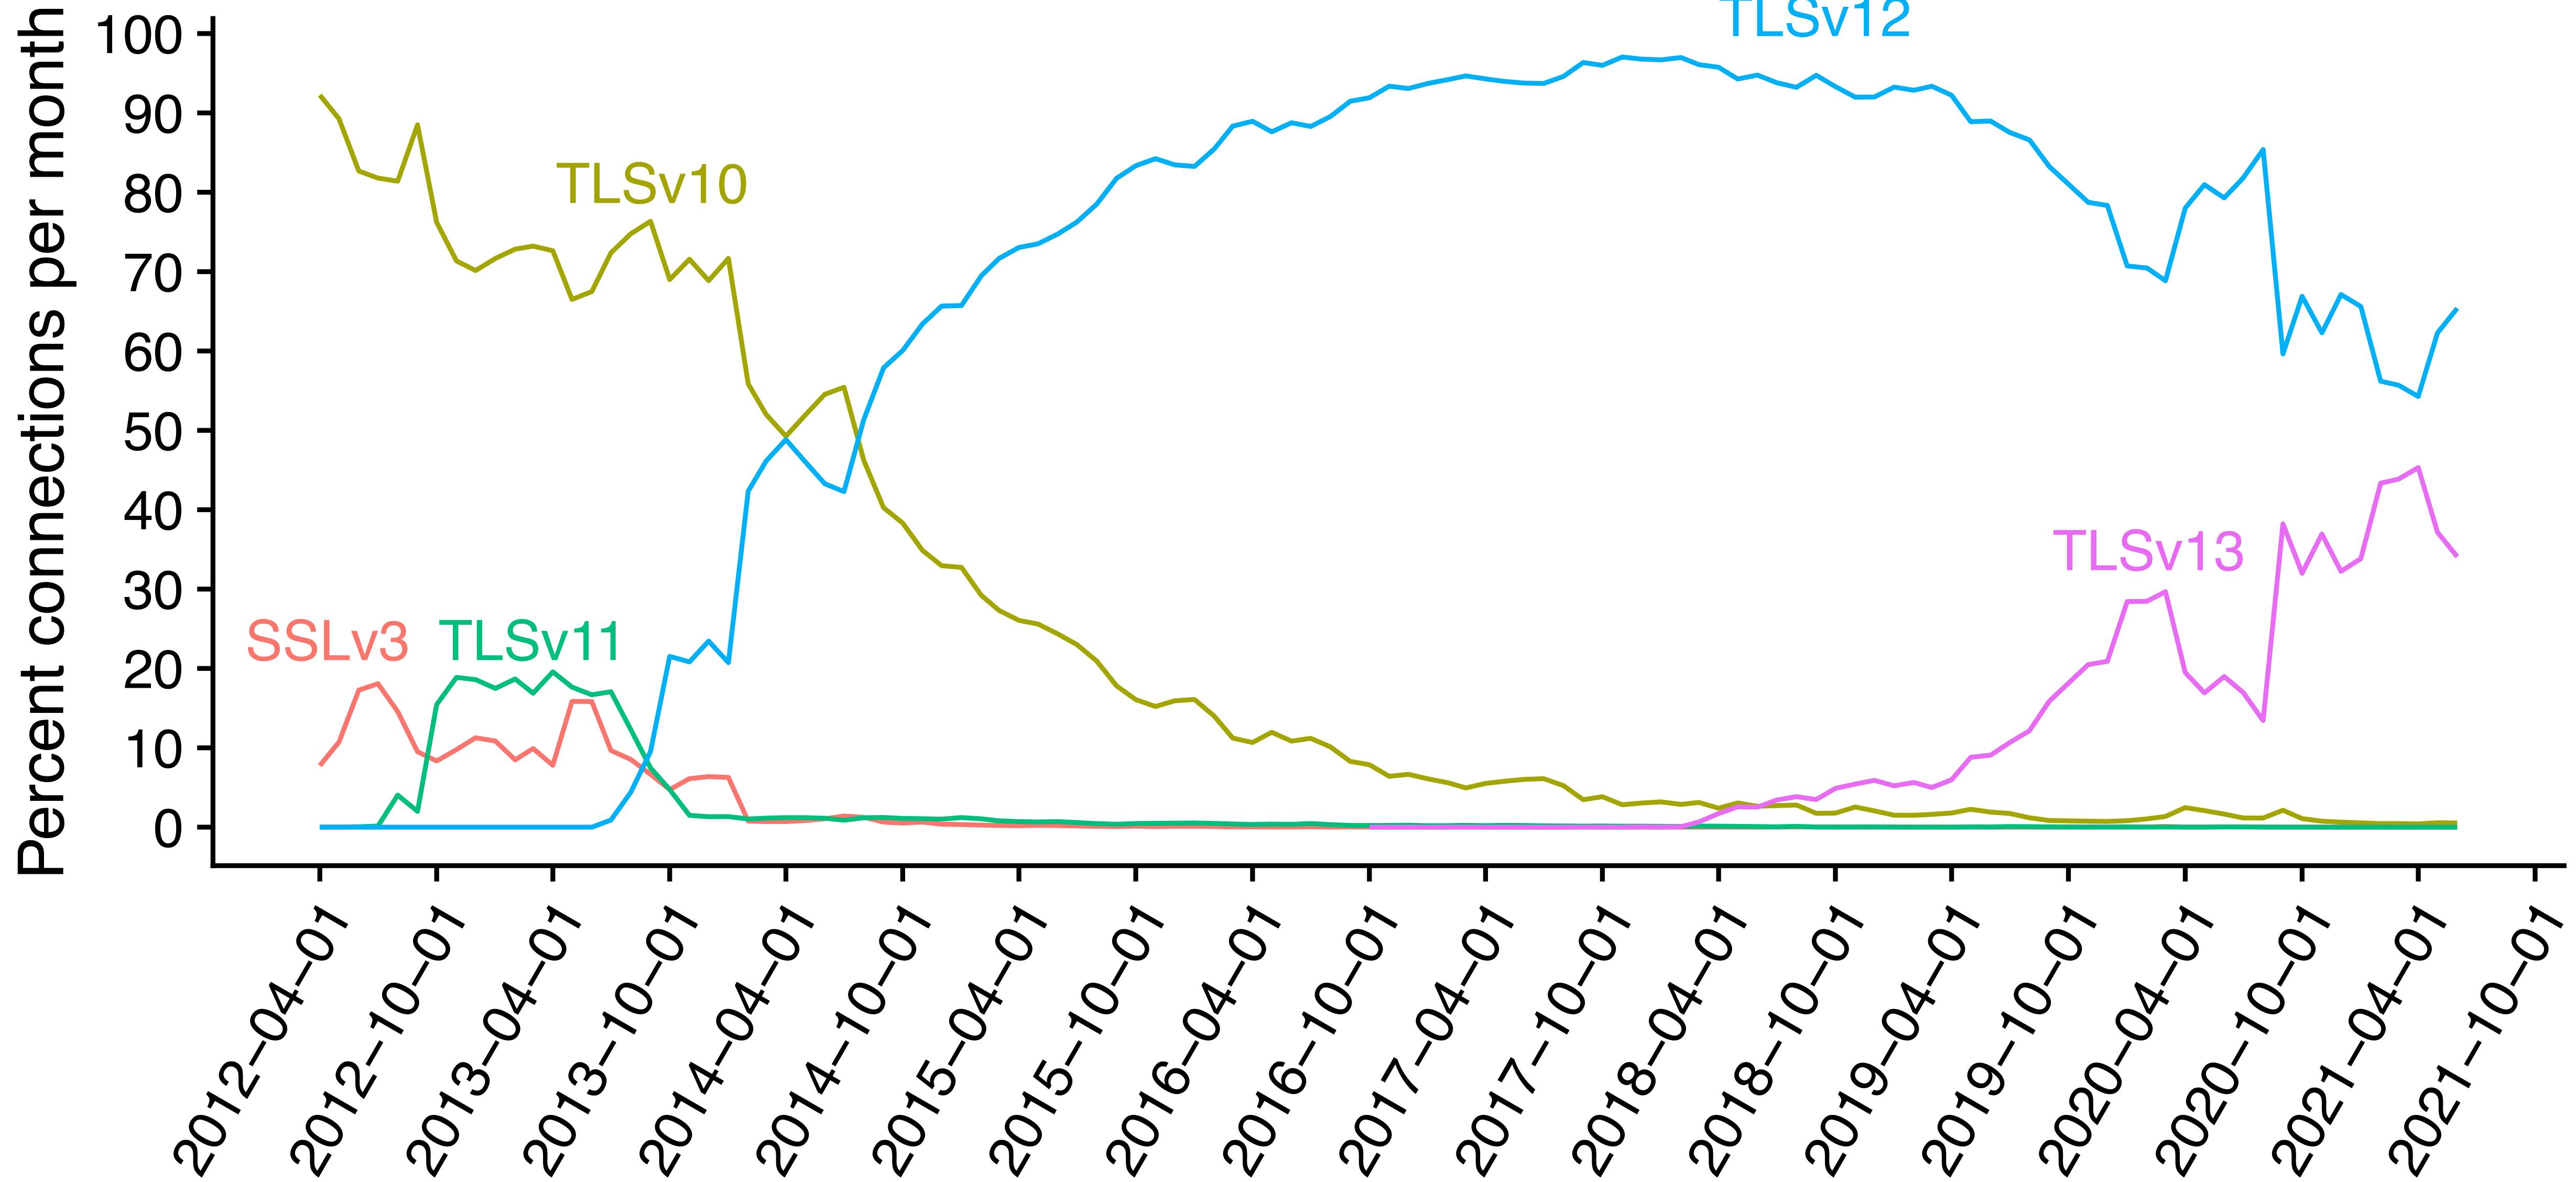


236 Million certificates

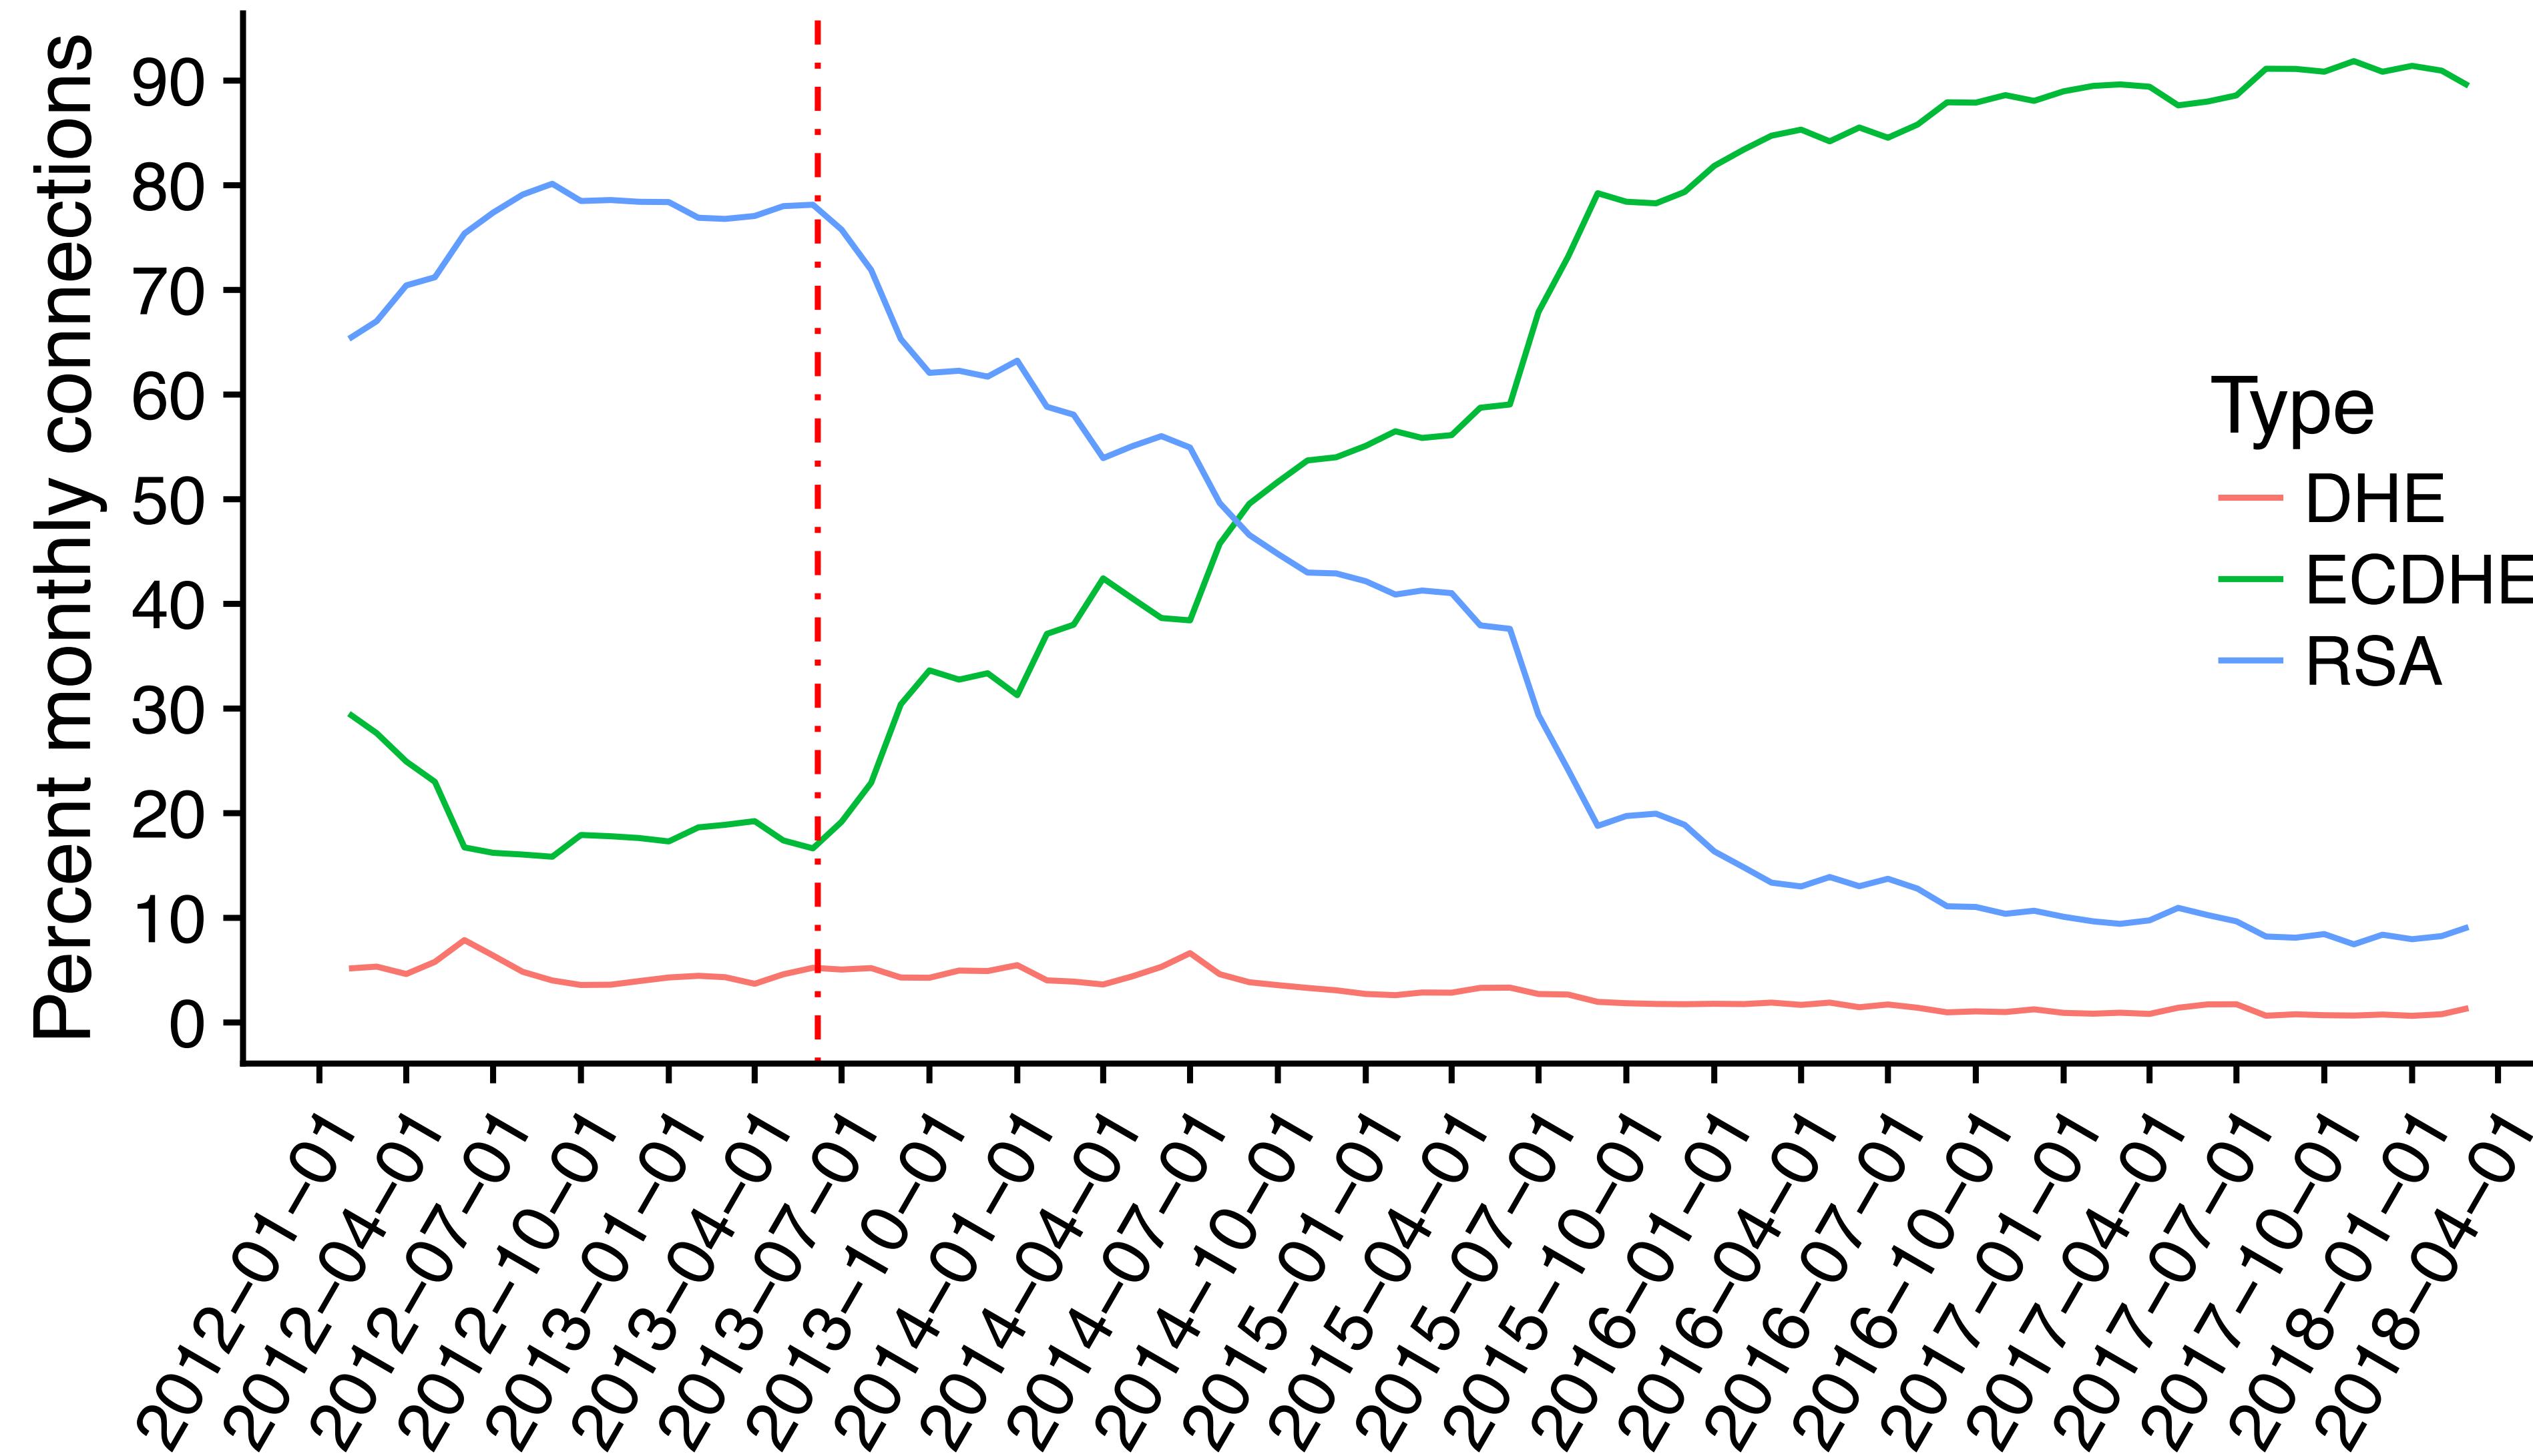
443 Billion connections



TLS Versions



Negotiated RSA vs forward secret



What is a certificate?

Version	3 (0x2)
Serial	0b:52:36:05:c5:be:20...
Signature Algorithm	sha256WithRSAEncryption
Issuer	C = US, O = Google Trust Services LLC, [...]
Valid NotBefore	Feb 28 02:21:45 2022 GMT
Valid NotAfter	May 23 02:21:44 2022 GMT
Subject	CN = *.google.com
SPKI (Key Information)	[...]
Extensions	[...]

What's weird about this one?

Version	3 (0x2)
Serial	0b:52:36:05:[..]
Signature Algorithm	sha256WithRSAEncryption
Issuer	CN=idrac-[...], O=Dell Inc.
Valid NotBefore	Feb 28 14:52:00 2016 GMT
Valid NotAfter	Feb 29 14:52:00 2026 GMT
Subject	CN=idrac-[...], O=Dell Inc.
SPKI (Key Information)	[...]
Extensions	[...]

What's weird about this one?

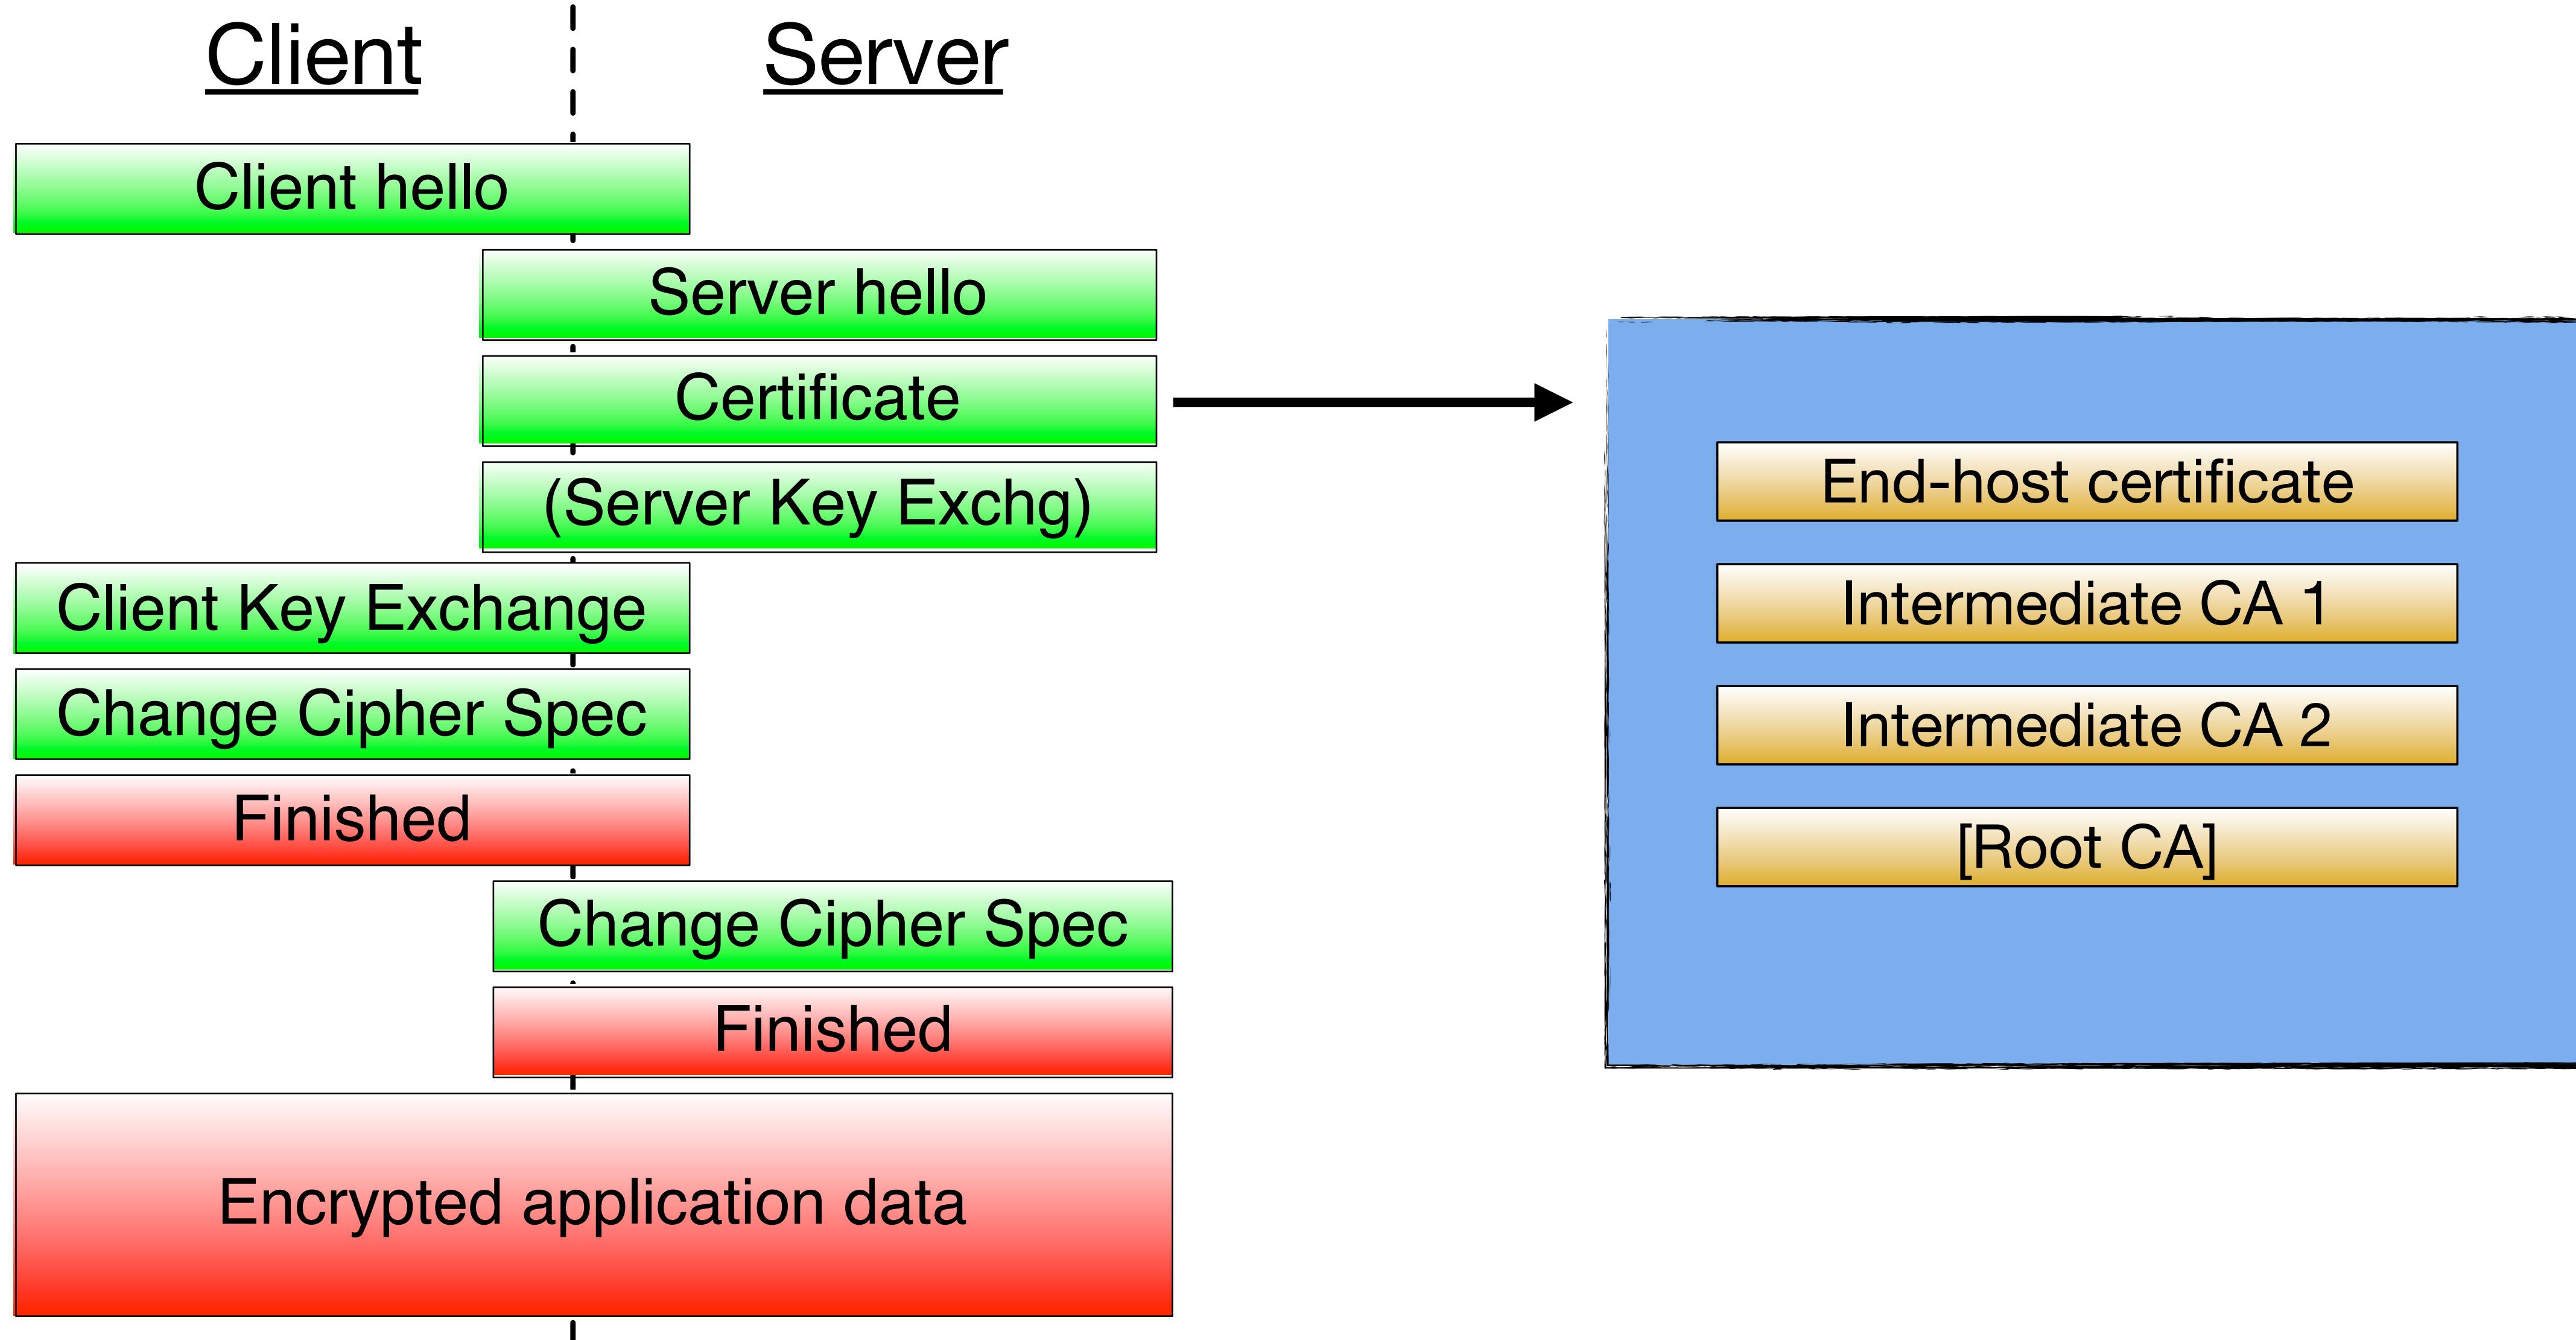
Version	3 (0x2)
Serial	0b:52:36:05:[..]
Signature Algorithm	sha256WithRSAEncryption
Issuer	CN=idrac-[...], O=Dell Inc.
Valid NotBefore	Feb 28 14:52:00 2016 GMT
Valid NotAfter	Feb 29 14:52:00 2026 GMT
Subject	CN=idrac-[...], O=Dell Inc.
SPKI (Key Information)	[...]
Extensions	[...]

One last one

Version	4 (0x3)
Serial	32112[.]
Signature Algorithm	sha256WithRSAEncryption
Issuer	O = Acronis INC, [...] CN = Datacenter CA
Valid NotBefore	Jul 4 06:55:07 2019 GMT
Valid NotAfter	Jul 3 06:55:07 2022 GMT
Subject	CN = Registration Server CA, [...]
SPKI (Key Information)	[...]
Extensions	[...]

One last one

Version	4 (0x3)
Serial	32112[...]
Signature Algorithm	sha256WithRSAEncryption
Issuer	O = Acronis INC, [...] CN = Datacenter CA
Valid NotBefore	Jul 4 06:55:07 2019 GMT
Valid NotAfter	Jul 3 06:55:07 2022 GMT
Subject	CN = Registration Server CA, [...]
SPKI (Key Information)	[...]
Extensions	[...]



Client

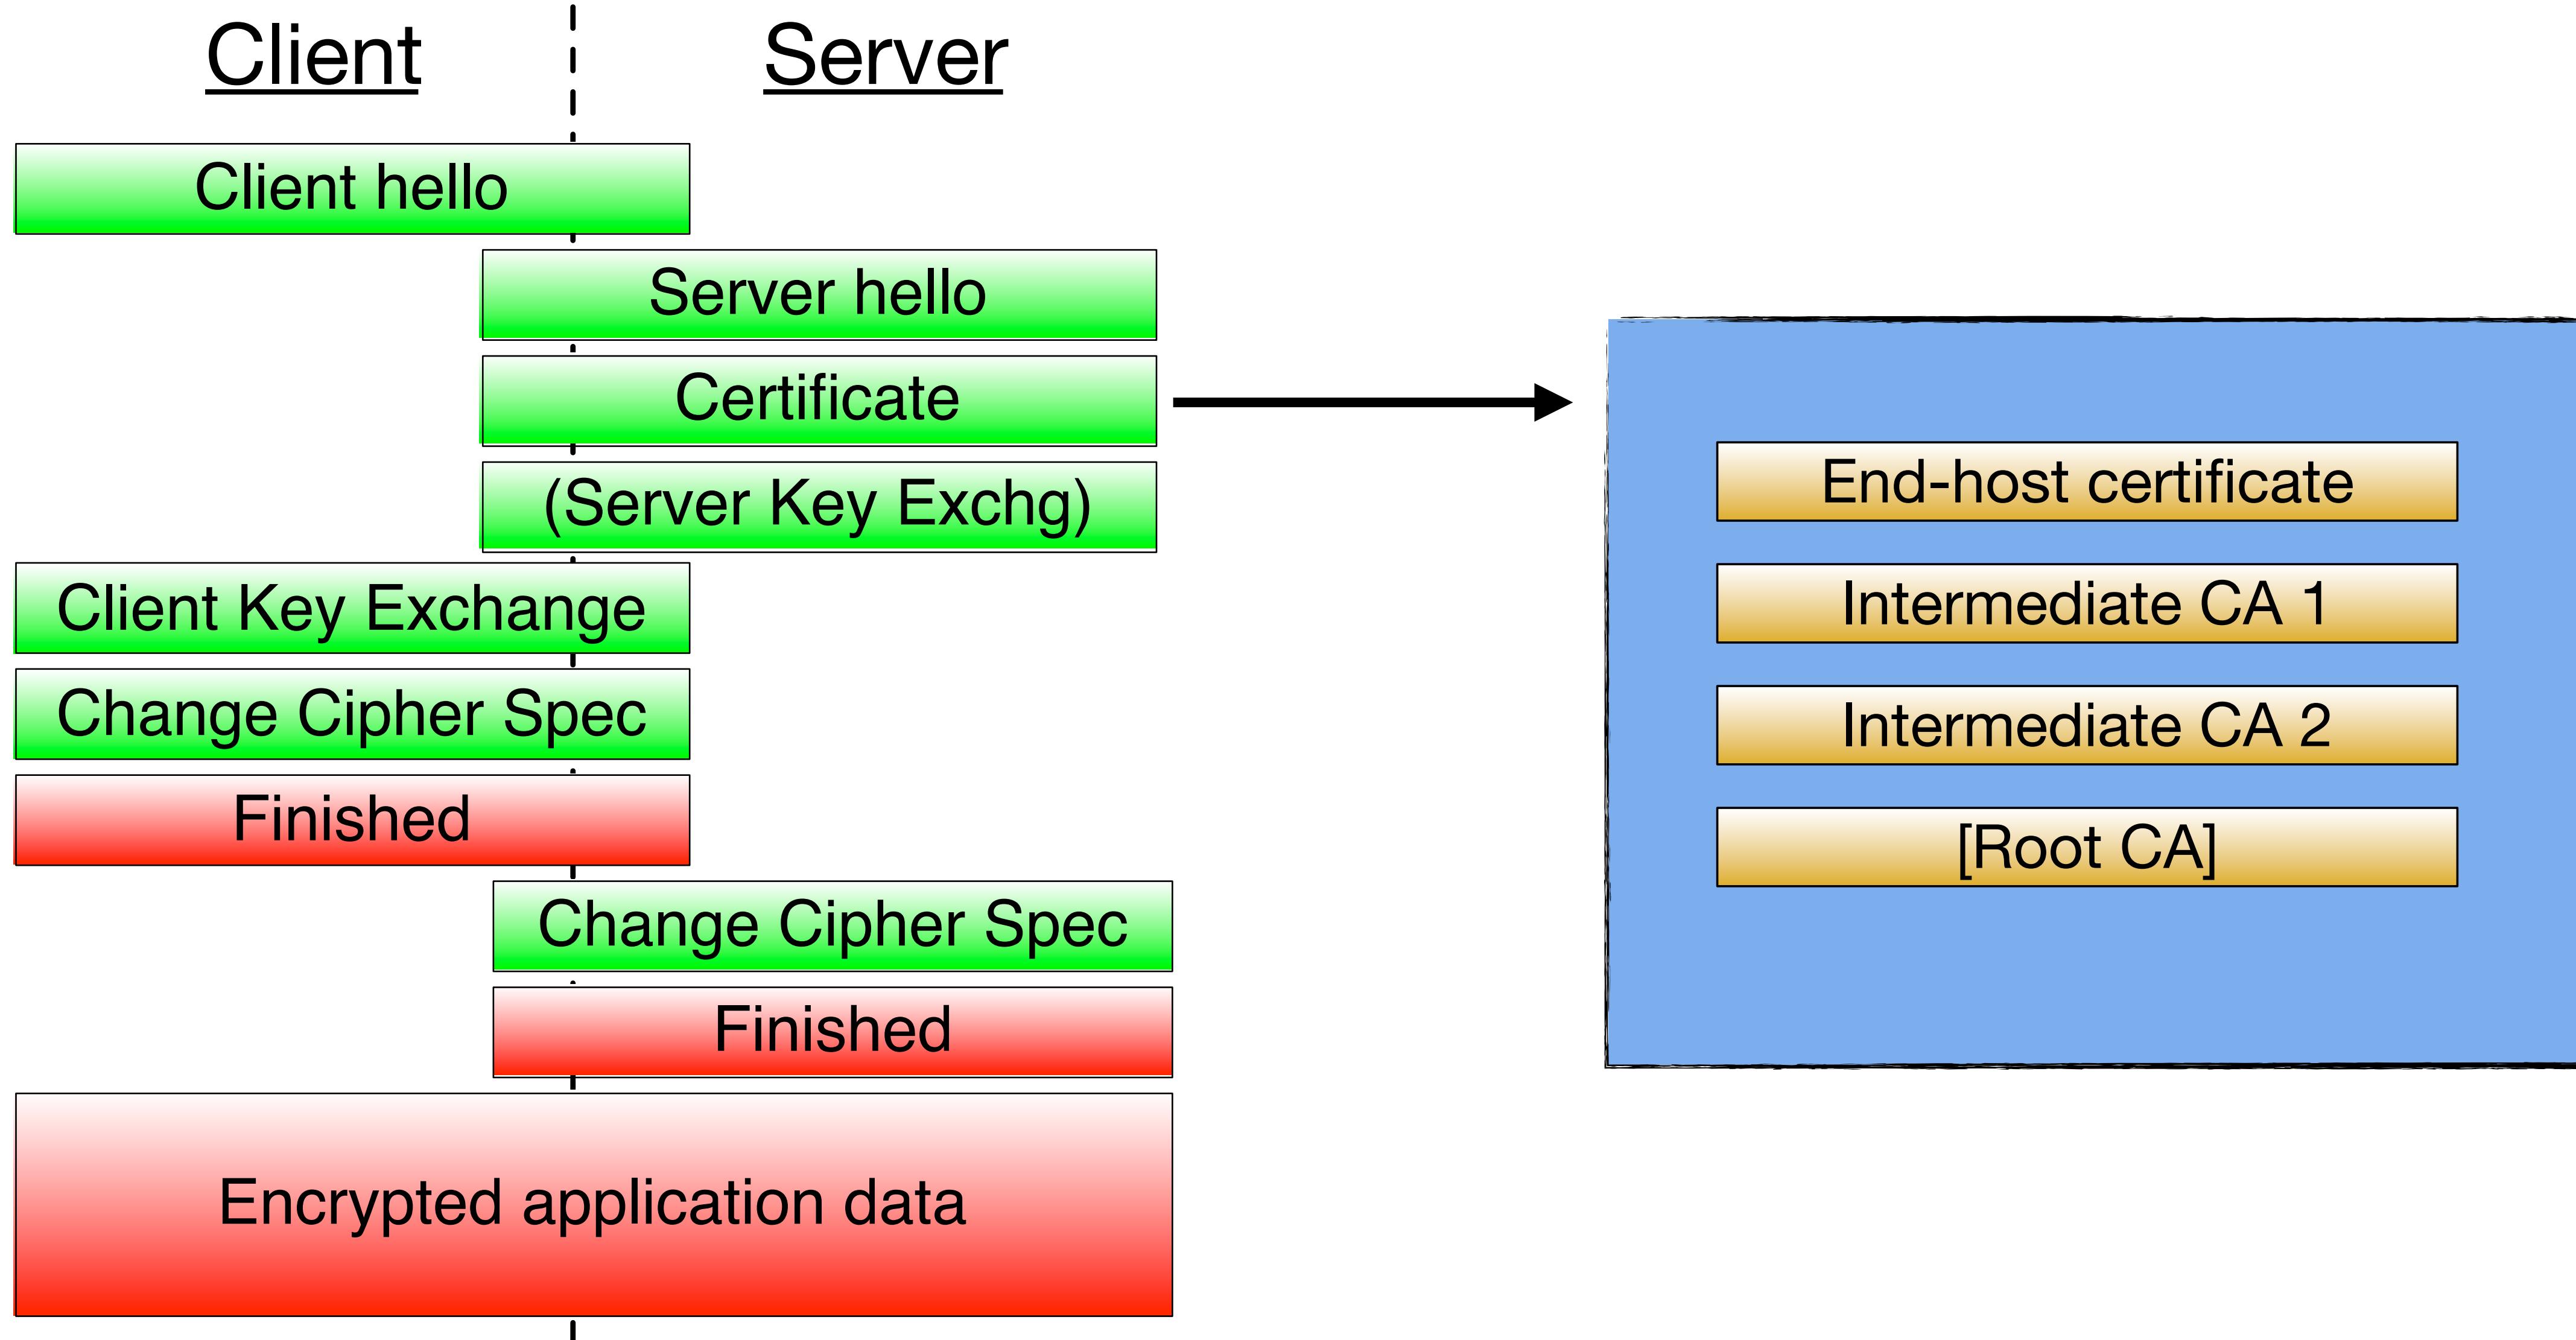
Server

Client hello

certificate_list

This is a sequence (chain) of certificates. The sender's certificate MUST come first in the list. Each following certificate MUST directly certify the one preceding it. Because certificate validation requires that root keys be distributed independently, the self-signed certificate that specifies the root certificate authority MAY be omitted from the chain, under the assumption that the remote end must already possess it in order to validate it in any case.

Encrypted application data



In Practice...

End-host certificate

No intermediates - common on Microsoft properties

End-host certificate

Old End-host certificate

Old certificates, or test certificates in the chain

Intermediates

Also - data after the certificate. Typically ignored during parsing.

```
Certificate:  
Data:  
Version: 3 (0x2)  
Serial Number:  
36:9a:c7:3d:67:06:3a:a2:75:83:0d:fc:66:84:1c:1e  
Signature Algorithm: sha256WithRSAEncryption  
Issuer: C=US, O=Symantec Corporation, OU=Symantec Trust Network, CN=Symantec Class 3 Secure Server CA - G4  
Validity  
Not Before: May 30 00:00:00 2016 GMT  
Not After : May 30 00:00:00 2018 GMT  
Subject: C=US, ST=Washington, L=Seattle, O=Amazon.com, Inc., CN=*.cloudfront.net  
X509v3 extensions:  
X509v3 Subject Alternative Name:  
DNS:cloudfront.net, DNS:*.cloudfront.net  
X509v3 Basic Constraints:  
CA:FALSE  
Authority Information Access:  
OCSP - URI:http://ss.symcd.com  
CA Issuers - URI:http://ss.symcb.com/ss.crt  
CT Precertificate SCTs:  
..Random string goes here
```

```
Certificate:  
Data:  
Version: 3 (0x2)  
Serial Number:  
36:9a:c7:3d:67:06:3a:a2:75:83:0d:fc:66:84:1c:1e  
Signature Algorithm: sha256WithRSAEncryption  
Issuer: C=US, O=Symantec Corporation, OU=Symantec Trust Network, CN=Symantec Class 3 Secure Server CA - G4  
Validity  
Not Before: May 30 00:00:00 2016 GMT  
Not After : May 30 00:00:00 2018 GMT  
Subject: C=US, ST=Washington, L=Seattle, O=Amazon.com, Inc., CN=*.cloudfront.net  
X509v3 extensions:  
X509v3 Subject Alternative Name:  
DNS:cloudfront.net, DNS:*.cloudfront.net  
X509v3 Basic Constraints:  
CA:FALSE  
Authority Information Access:  
OCSP - URI:http://ss.symcd.com  
CA Issuers - URI:http://ss.symcb.com/ss.crt  
CT Precertificate SCTs:  
..Random string goes here
```

853:d=5 hl=2 l= 3 prim: OBJECT	:X509v3 CRL Distribution Points
858:d=5 hl=2 l= 36 prim: OCTET STRING	[HEX DUMP]:30223020A01EA01C861A687474703A2F2F73732E73796D63622E636F6D2F73732E63726C
896:d=4 hl=2 l= 87 cons: SEQUENCE	
898:d=5 hl=2 l= 8 prim: OBJECT	:Authority Information Access
908:d=5 hl=2 l= 75 prim: OCTET STRING	[HEX DUMP]:3049301F06082B060105050730018613687474703A2F2F73732E73796D63642E636F6D30260
063622E636F6D2F73732E637274	
985:d=4 hl=2 l= 39 cons: SEQUENCE	
987:d=5 hl=2 l= 10 prim: OBJECT	:CT Precertificate SCTs
999:d=5 hl=2 l= 25 prim: OCTET STRING	[HEX DUMP]:0C1752616E646F6D20737472696E6720676F65732068657265

853:d=5 hl=2 l= 3 prim: OBJECT	:X509v3 CRL Distribution Points
858:d=5 hl=2 l= 36 prim: OCTET STRING	[HEX DUMP]:30223020A01EA01C861A687474703A2F2F73732E73796D63622E636F6D2F73732E63726C
896:d=4 hl=2 l= 87 cons: SEQUENCE	
898:d=5 hl=2 l= 8 prim: OBJECT	:Authority Information Access
908:d=5 hl=2 l= 75 prim: OCTET STRING	[HEX DUMP]:3049301F06082B060105050730018613687474703A2F2F73732E73796D63642E636F6D30260
063622E636F6D2F73732E637274	
985:d=4 hl=2 l= 39 cons: SEQUENCE	
987:d=5 hl=2 l= 10 prim: OBJECT	:CT Precertificate SCTs
999:d=5 hl=2 l= 25 prim: OCTET STRING	[HEX DUMP]:0C1752616E646F6D20737472696E6720676F65732068657265

```
$ openssl asn1parse -in invalidsct.crt -inform der -strparse 999
 0:d=0  hl=2 l= 23 prim: UTF8STRING      :Random string goes here
```

Normal SCT

CT Precertificate SCTs:

Signed Certificate Timestamp:

Version : v1(0)

Log ID : DD:EB:1D:2B:7A:0D:4F:A6:20:8B:81:AD:81:68:70:7E:
2E:8E:9D:01:D5:5C:88:8D:3D:11:C4:CD:B6:EC:BE:CC

Timestamp : Aug 17 17:25:11.747 2016 GMT

Extensions: none

Signature : ecdsa-with-SHA256

30:46:02:21:00:B9:6C:2B:9A:D5:C8:70:EC:CD:2E:17:
E6:69:5E:C0:51:47:24:D5:DE:37:CF:10:54:84:A7:D6:
FD:6B:A4:A6:31:02:21:00:ED:0C:E0:49:63:60:D7:26:
DD:DD:06:B4:80:D6:42:FC:F4:C5:74:70:C5:4F:4D:8D:
9F:41:61:91:BB:B1:73:86

Signed Certificate Timestamp:

Version : v1(0)

Log ID : A4:B9:09:90:B4:18:58:14:87:BB:13:A2:CC:67:70:0A:
3C:35:98:04:F9:1B:DF:B8:E3:77:CD:0E:C8:0D:DC:10

Timestamp : Aug 17 17:25:11.810 2016 GMT

Extensions: none

Signature : ecdsa-with-SHA256

30:45:02:21:00:C4:A9:7D:4B:93:C1:57:BB:AF:39:01:
D9:5B:CB:01:35:44:97:7A:9B:E9:FD:A2:F7:15:CA:F2:
16:4B:88:5E:AC:02:20:10:9D:1E:54:8D:3A:C1:20:65:
A9:25:BE:8F:00:8E:26:26:2D:D8:E7:BA:AE:48:84:19:
35:86:0D:B8:EC:B3:D4

```
Certificate:  
Data:  
Version: 3 (0x2)  
Serial Number:  
36:9a:c7:3d:67:06:3a:a2:75:83:0d:fc:66:84:1c:1e  
Signature Algorithm: sha256WithRSAEncryption  
Issuer: C=US, O=Symantec Corporation, OU=Symantec Trust Network, CN=Symantec Class 3 Secure Server CA - G4  
Validity  
Not Before: May 30 00:00:00 2016 GMT  
Not After : May 30 00:00:00 2018 GMT  
Subject: C=US, ST=Washington, L=Seattle, O=Amazon.com, Inc., CN=*.cloudfront.net  
X509v3 extensions:  
X509v3 Subject Alternative Name:  
DNS:cloudfront.net, DNS:*.cloudfront.net  
X509v3 Basic Constraints:  
CA:FALSE  
Authority Information Access:  
OCSP - URI:http://ss.symcd.com  
CA Issuers - URI:http://ss.symcb.com/ss.crt  
CT Precertificate SCTs:  
..Random string goes here
```

A small list of things we have seen:

- Invalid hostnames
- Invalid certificate versions...
- Extra certificates
- Unordered certificates
- Missing certificates
- Extra bytes after certificates
- Data that is not certificates
- Autogenerated certificates
- Certificates with unusual choices in algorithms
- Fake certificates

Takeaways

- Check how your measurement infrastructure deals with outliers/edge behavior
- There will be something that you did not think of
- If something seems weird - dig into it.
It is fun - and there is a chance that you will find something interesting.
- When writing papers, figuring out all the edge-cases takes quite a while.
This stings, especially when you are rejected - lots of repeated work

Consider capturing your traffic...

...even when you perform active measurements

- You get repeatability
- You can go back to the raw data and check for things you might have missed...
 - Unimplemented protocol features
 - Responses that you recorded wrong
- This works even for encrypted traffic - you can save your key material!

Thank you

This work would not have been possible without the work of a lot of other people. Thank you very much to Robin Sommer, Christian Kreibich, Seth Hall, Matthias Vallentin, Aashish Sharma, Vern Paxson, Ralph Holz, Jens Hiller, Platon Kotzias, Kenny Paterson - and many others that worked on this over the years.

I also want to thank our data providers - many contributed data for years, fixed problems for us, helped us diagnose problems, explained their network setups to us, etc.

And finally - thank you to the photographers of Unsplash.

