Jon Nistor

nistor@torix.ca

Boulevard of Broken Networks

Agenda

Introduction

Networks

Findings

What you can do

Introduction

Who am I?

Who is Ops?

Introduction

Why this presentation?

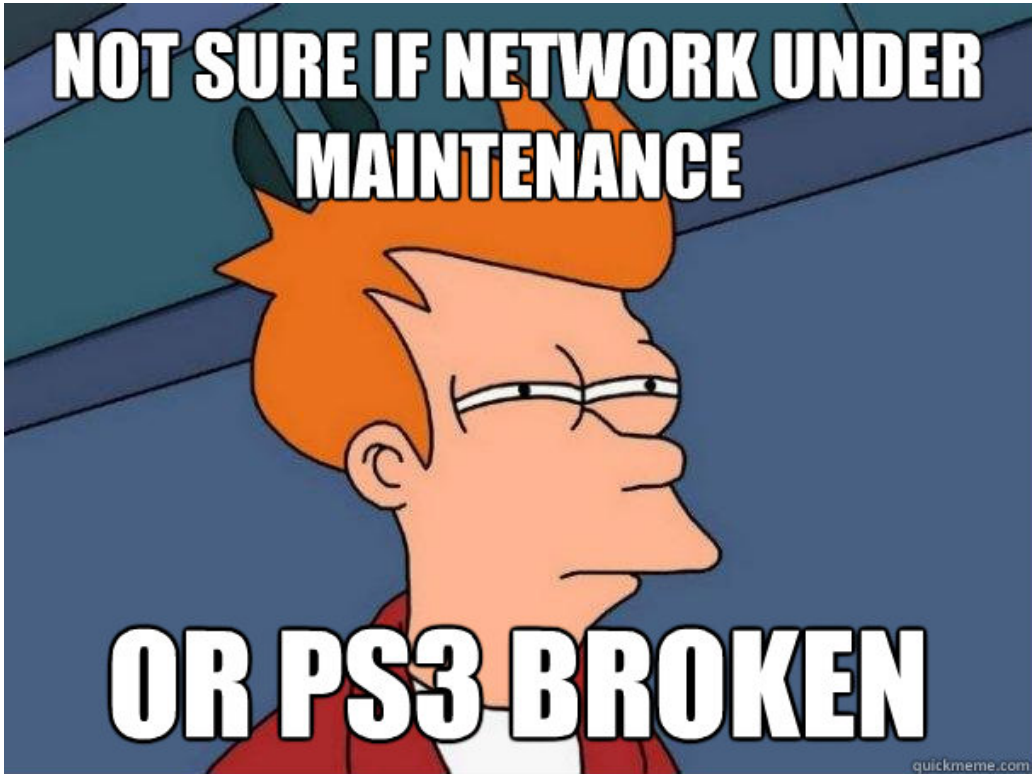

Introductions

What will we talk about?

How do we see all this?

We call in for backup

How do we see all this?

We are in the middle of the action

Attention to detail

How do we see all this?

Simple tools

- snoopdog
- tcpdump
- perl
- arpwatch
- snmp[trap]

Peer connections

Types of network layouts

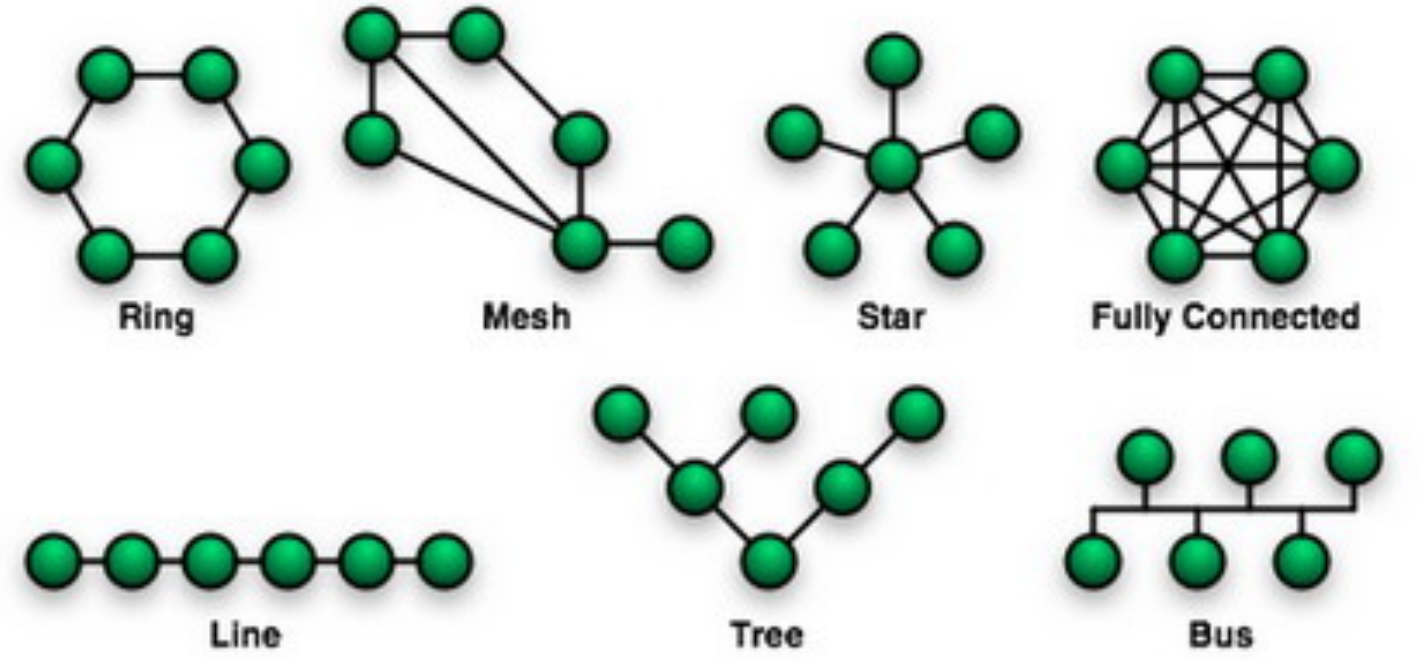

Network diagram 1

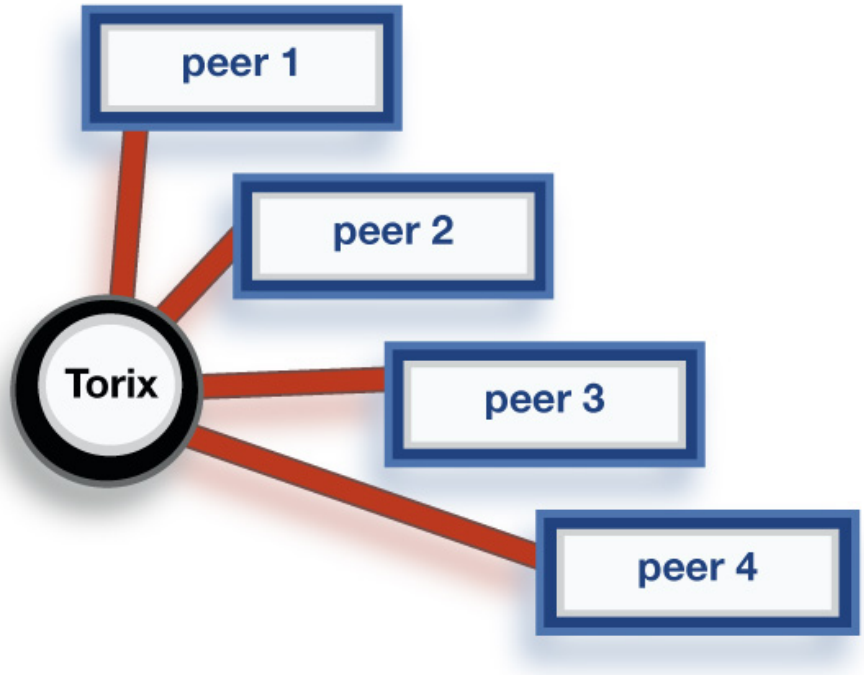

Network diagram 2

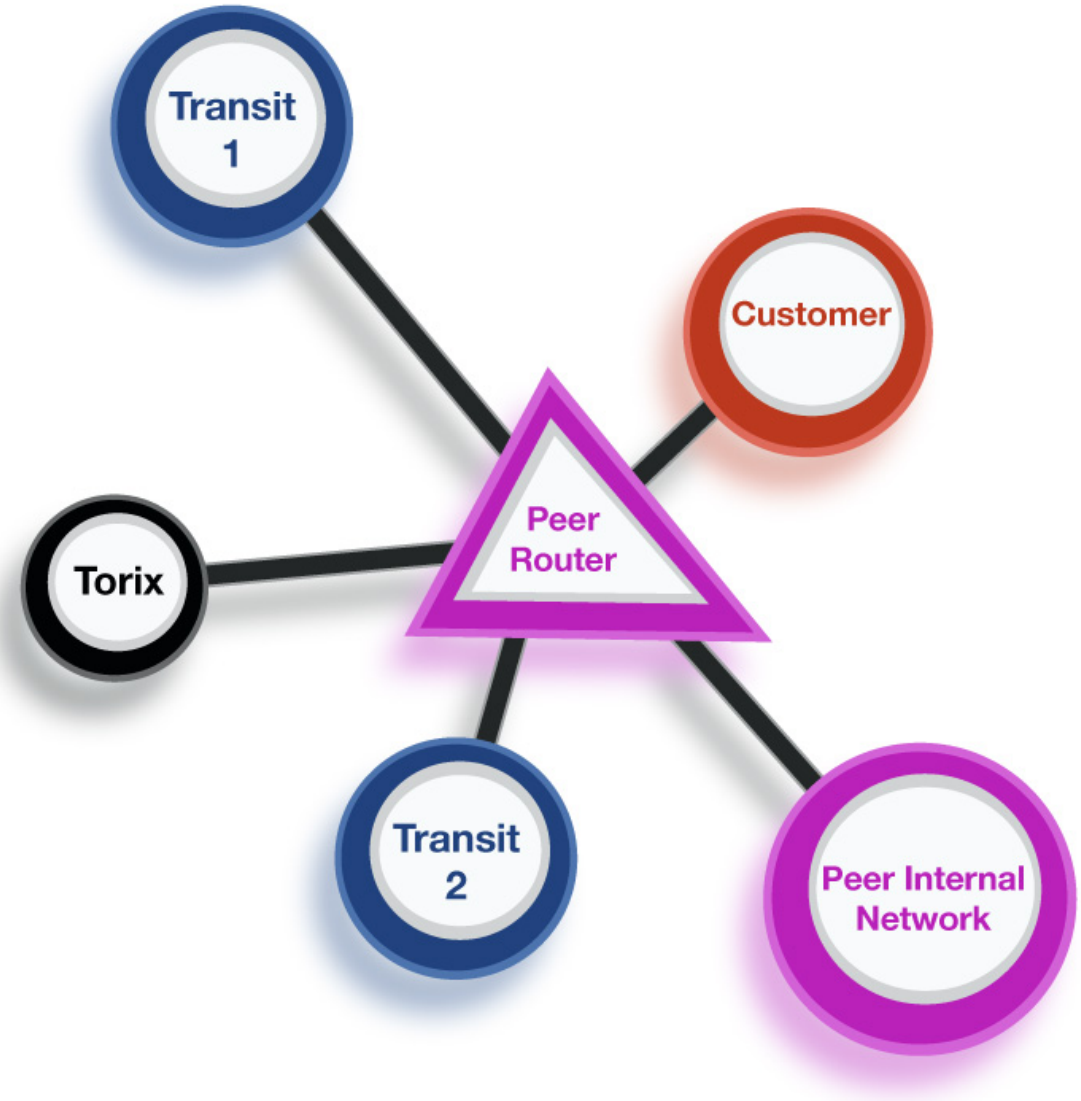

Network diagram 3

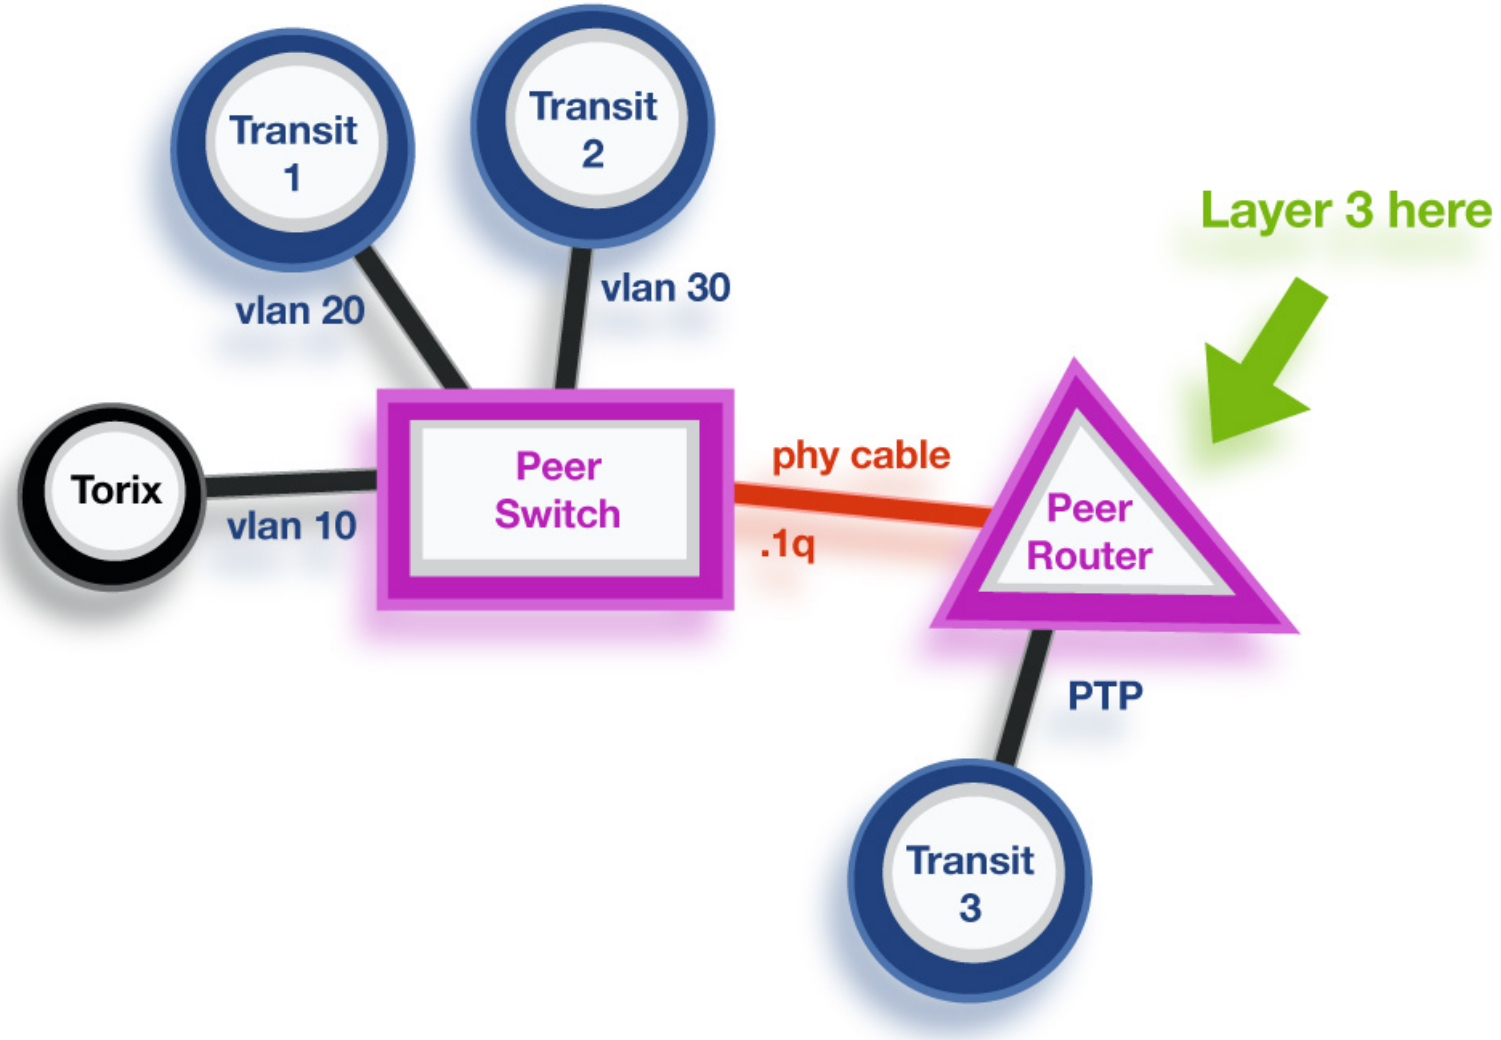

Network diagram 4

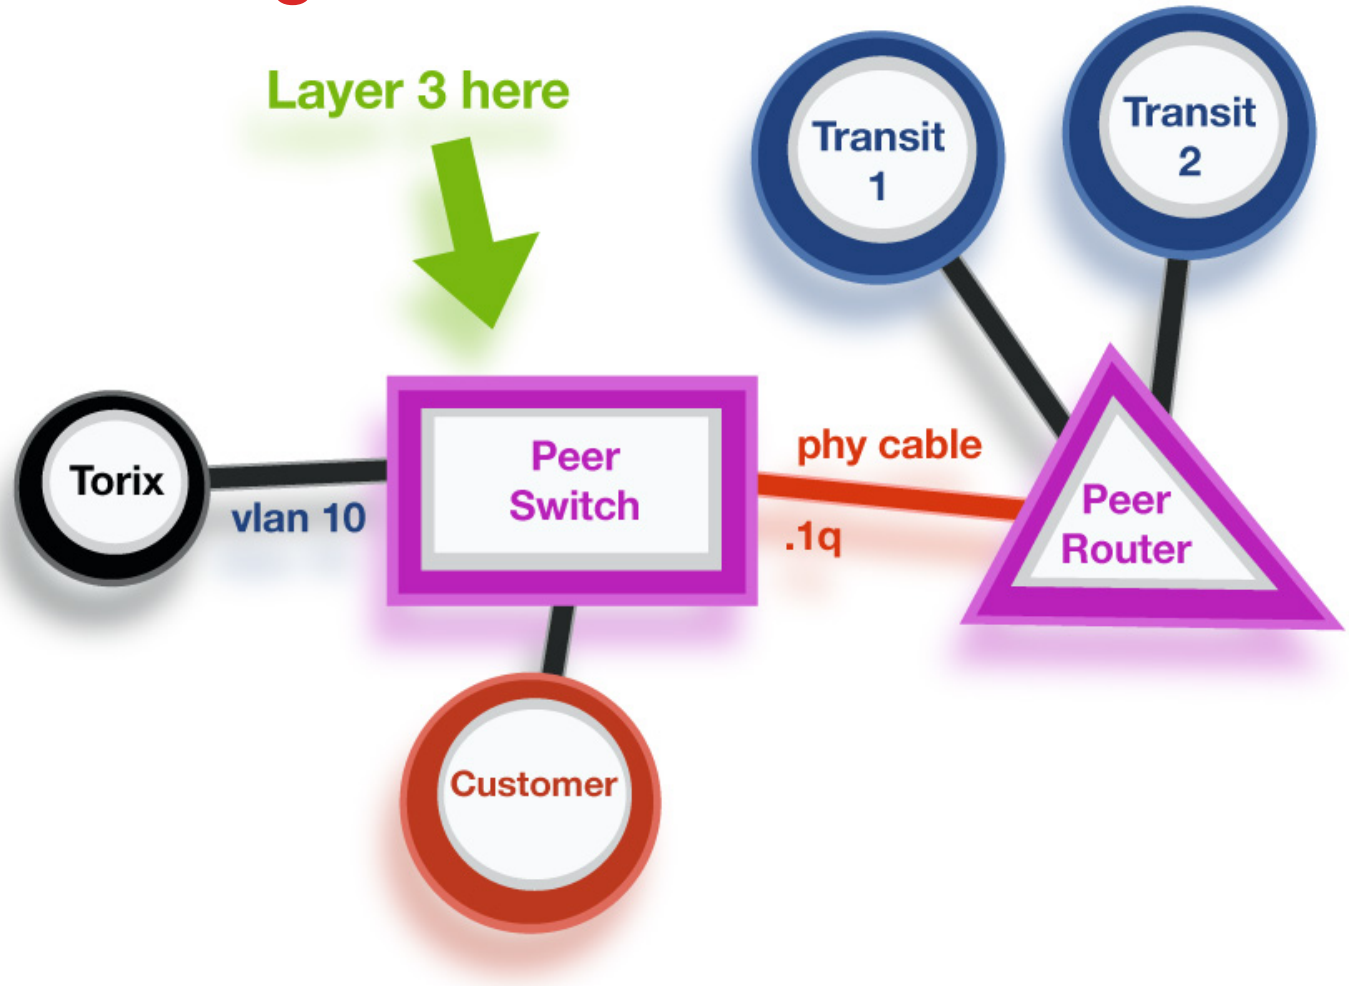

Network diagram 5

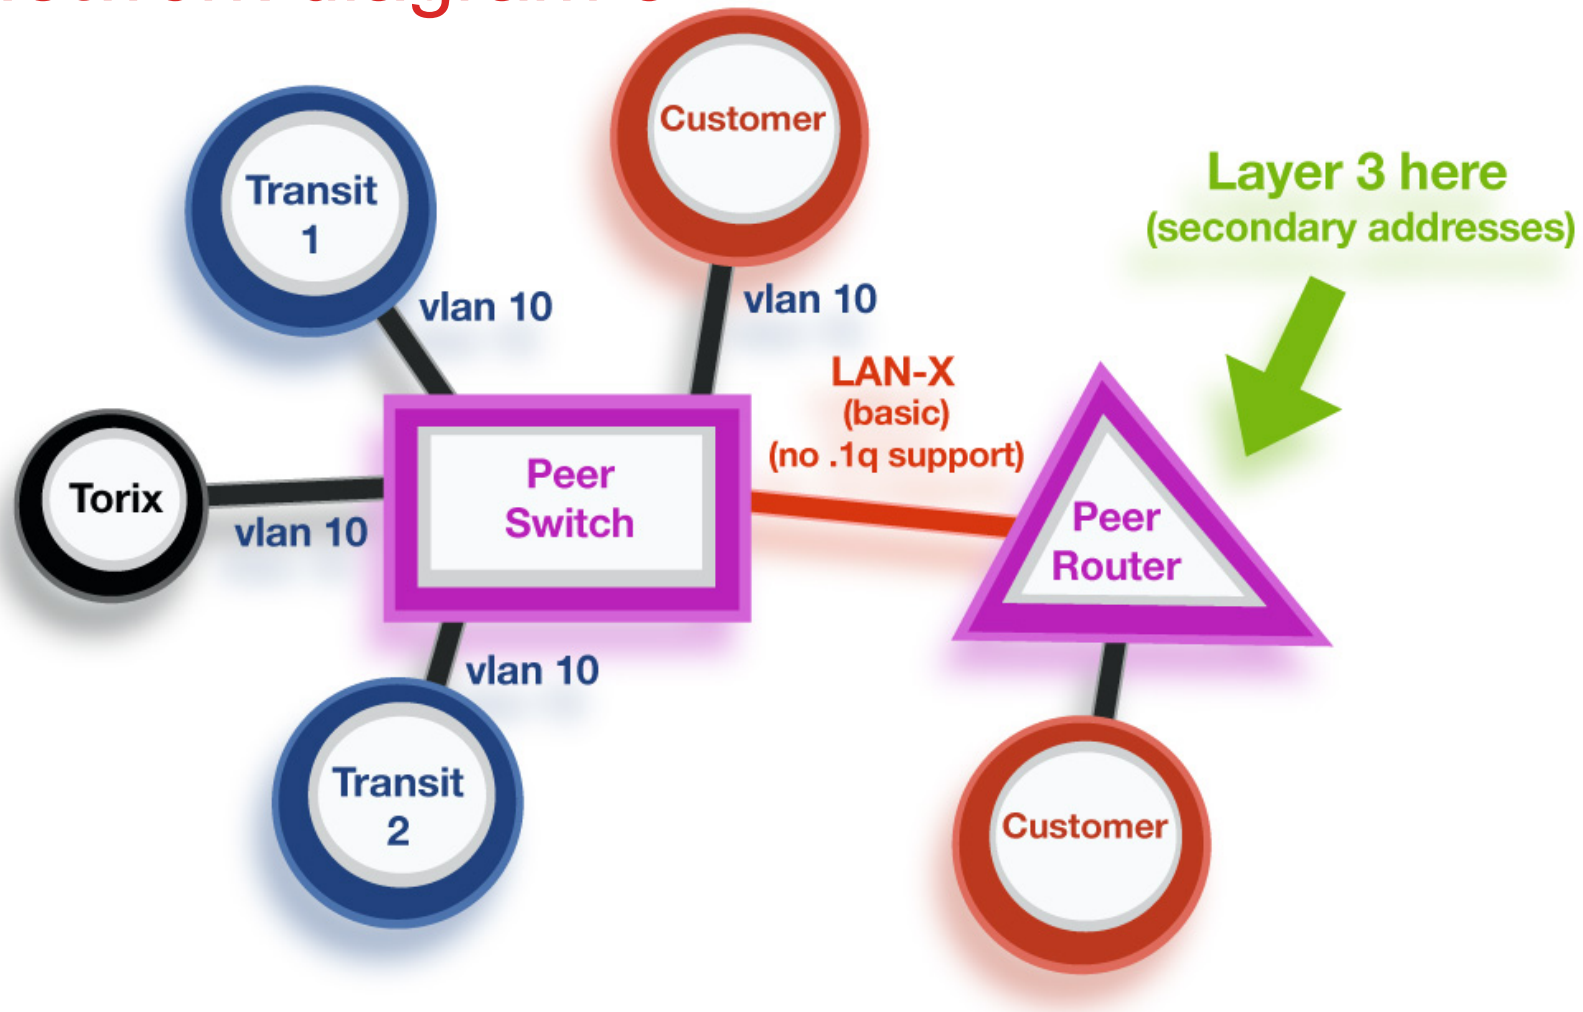

Network diagram 6

Network diagram 7

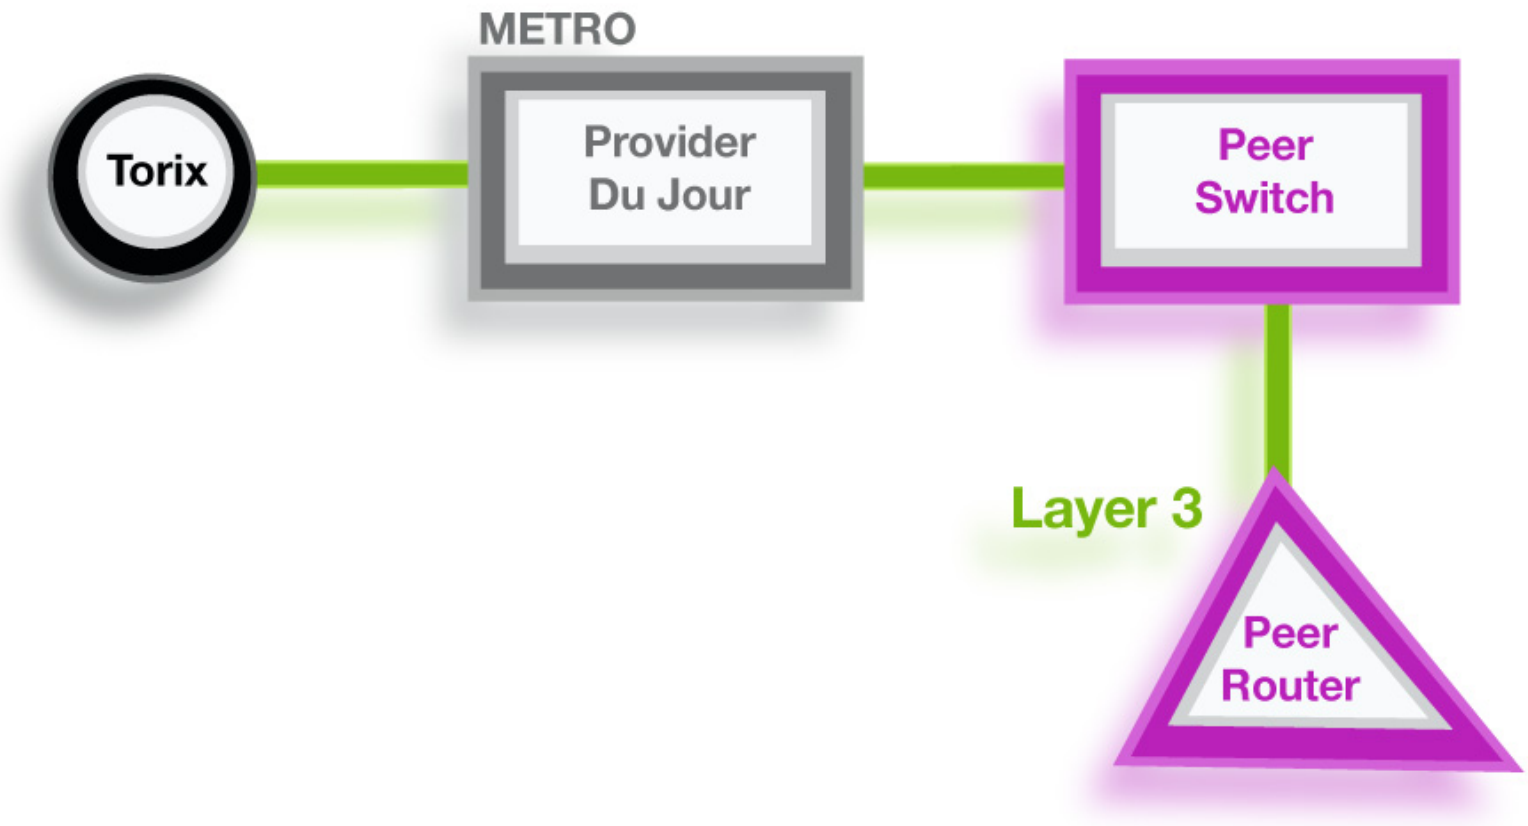

Network diagram 8

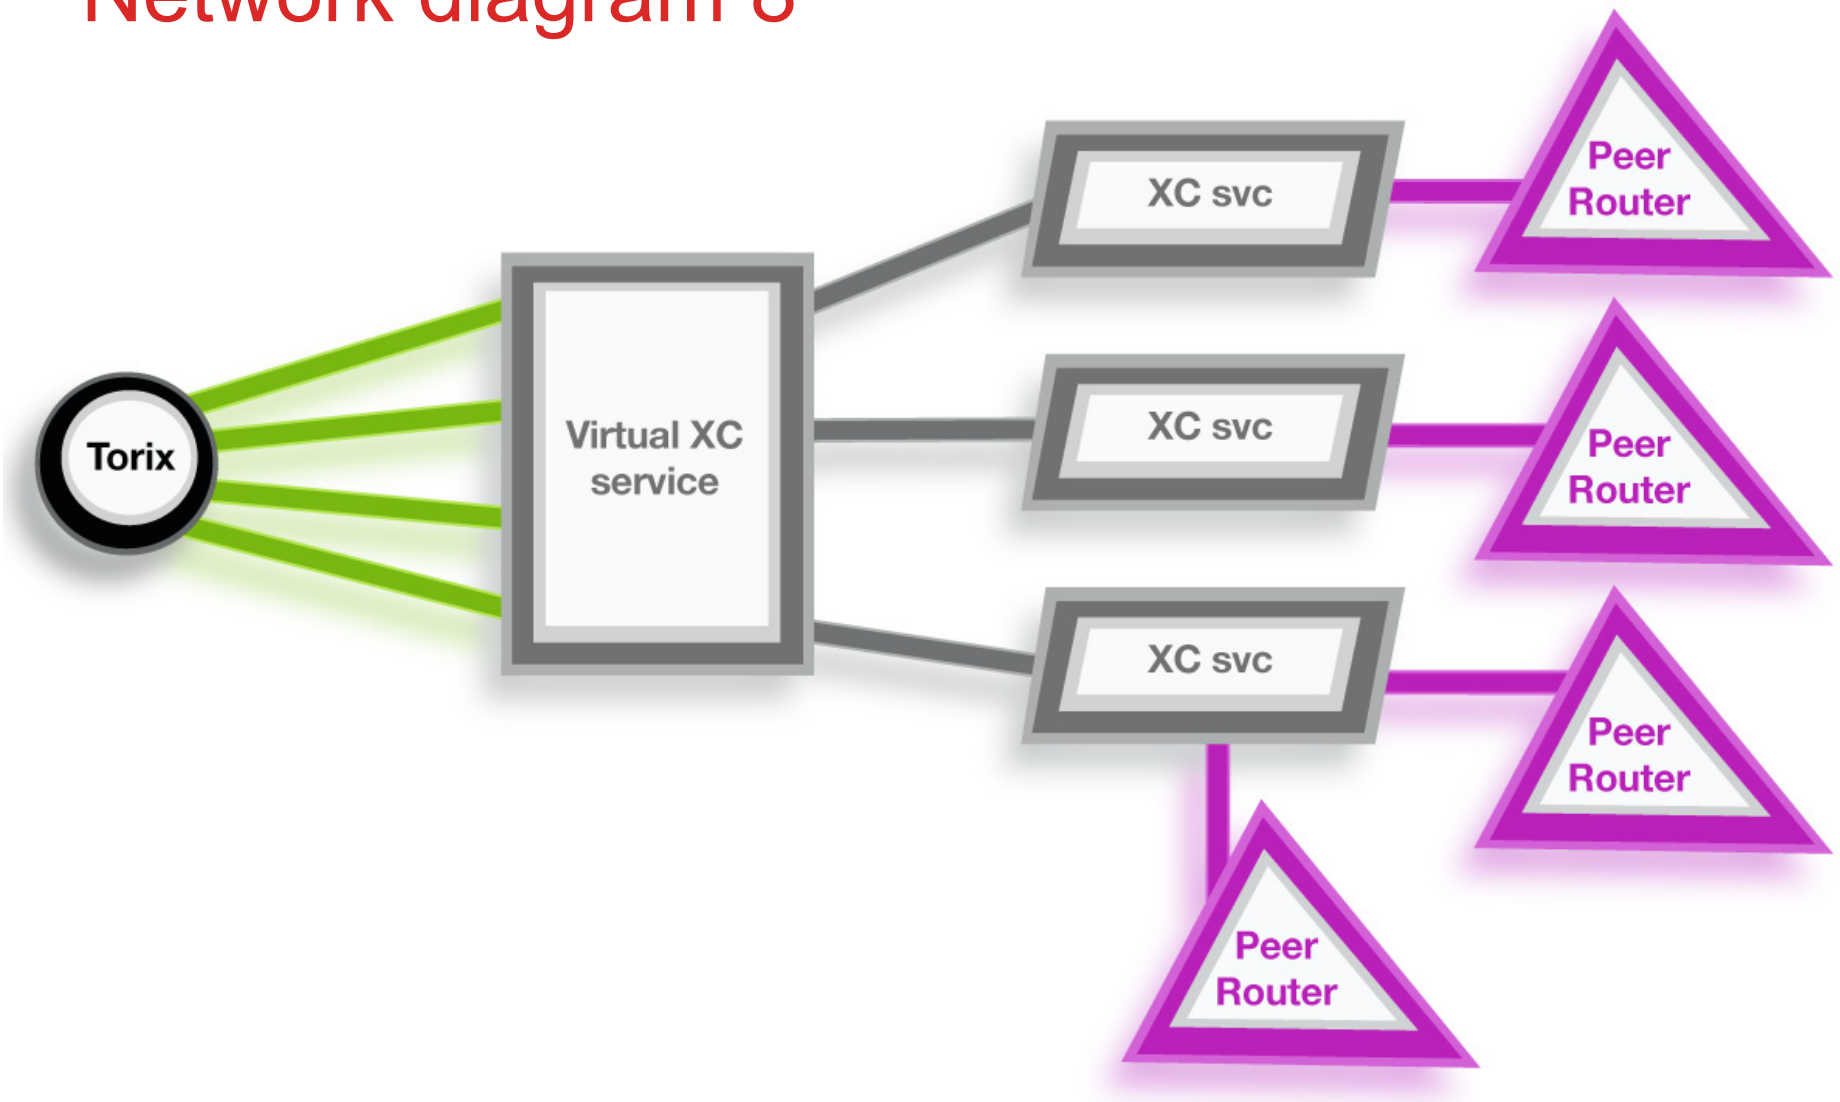

Physical

Where connectivity problems start

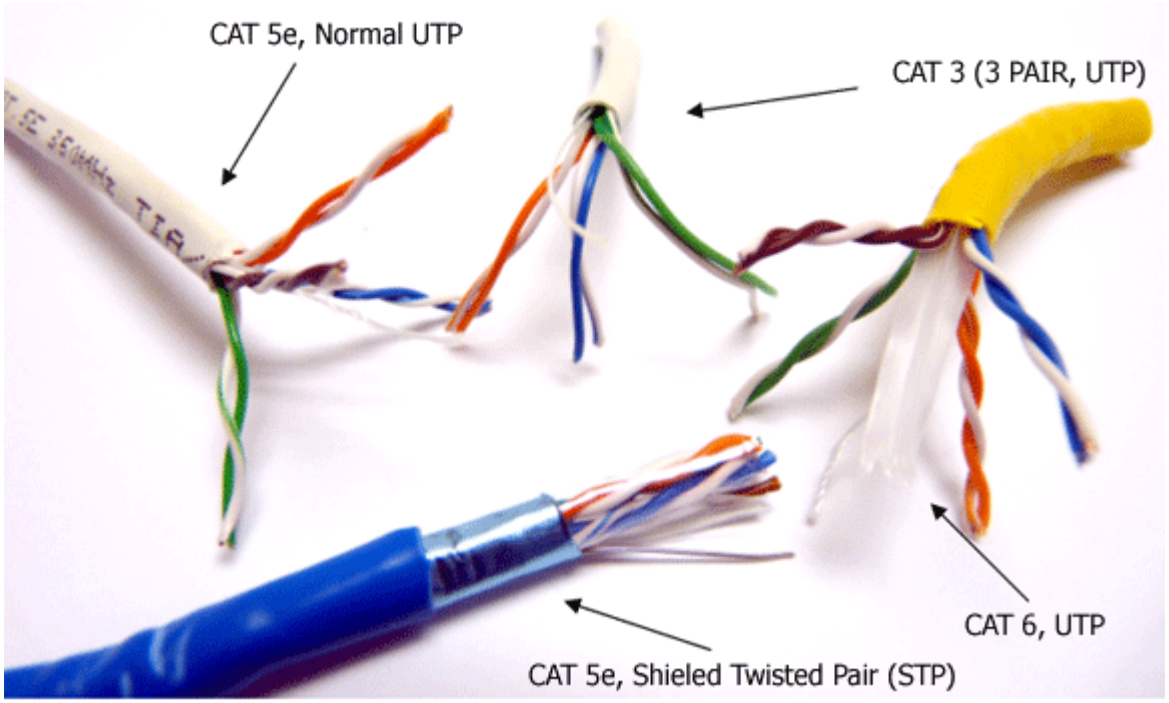

Physical: survey

Facebook survey:

"If you were to crimp 100 RJ45/Ethernet copper cables, how many would you get wrong for any reasons, either wrong pin-out, bad crimp, upside down ends, etc.?"

Polls: 83

Physical: survey

....survey says

cables	1-5~	62% people
	6-10~	22% …
	11-15~	5% …
	16-20~	11% …

Physical: survey

….survey had great suggestions

“I’d just put my monoprice express order in” (Paul)

“I’m with Paul, zero as my summer student would know his/her job depends on it.” (Rob)

Physical: optical

Splice or buy fiber?

Majority buy

Might clean

Scoping

OTDR

Physical: error tree

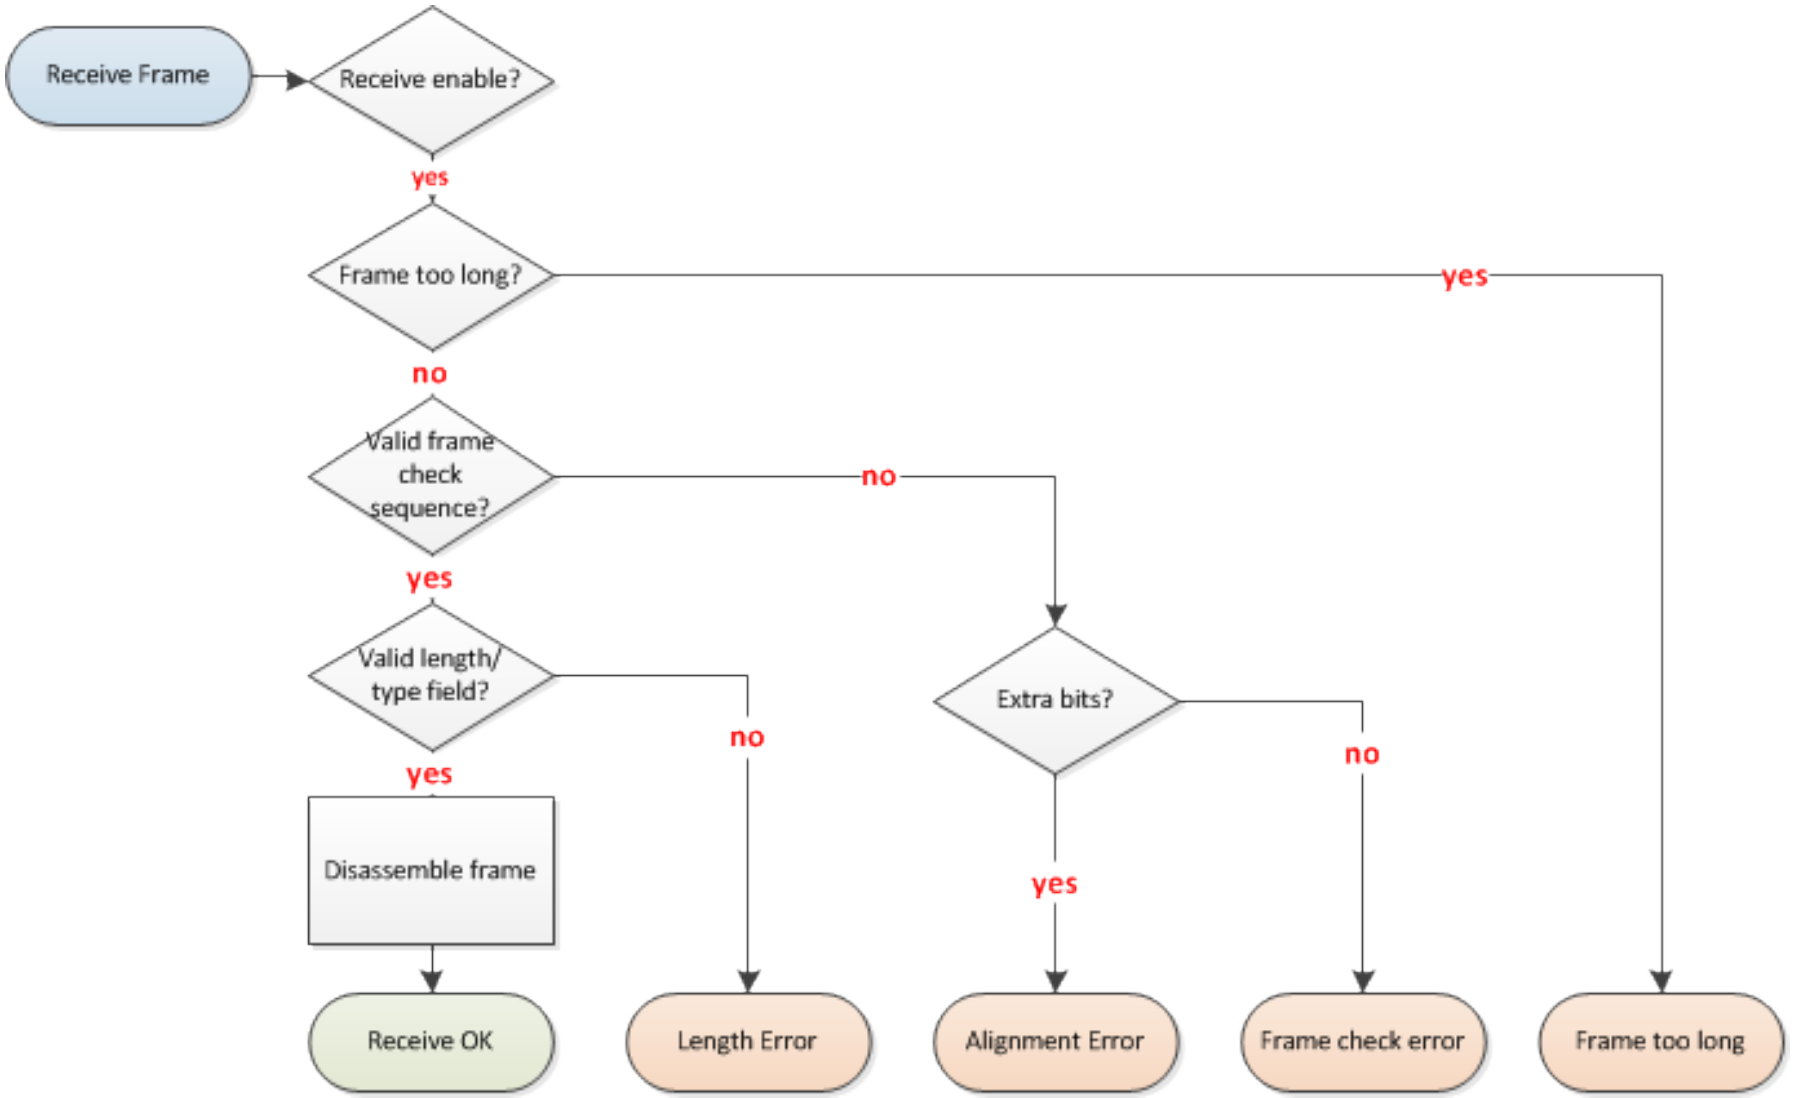

Physical: errors

Port	Align-Err	FCS-Err	Xmit-Err	Rcv-Err	UnderSize	OutDiscards
Eth1/27	0	0	0	5	0	0
Eth1/30	0	0	0	35	0	0
Eth2/21	4	415	0	449	0	0
Eth2/23	0	0	0	7	0	0
Eth2/24	0	87	0	97	0	0
Eth2/25	2	3756162	0	3795765	0	0
Eth2/30	0	216	0	247	0	0
Eth3/4	0	6102	0	6102	257831	0
Eth3/15	2	36241	0	36243	1479957	0
Eth3/18	0	0	0	121	0	0
Eth3/20	0	5	0	5	0	0
Eth3/26	0	0	0	16734	0	0
Eth3/30	0	0	0	8	0	253212352
Eth4/20	0	0	0	0	939	0
Eth7/6	0	0	0	975735	0	0

Port	Giants	SQETest-Err	Deferred-Tx	IntMacTx-Er	IntMacRx-Er	Symbol-Err
Eth3/15	0	--	0	0	0	2521768
Eth3/26	1350	--	0	0	0	0

Ethertypes and MACs

0x0800	IPv4
0x0806	ARP
0x6002	MOP/RC
0x86DD	IPv6*
0x8809	LACP*
0x88CC	LLDP

01:80:C2:00:00:00	IEEE 802.1D
FF:FF:FF:FF:FF:FF	Broadcast

STP: BPDUs

- Using single location for collection
- We always suggest peers turn off stp, why do you even need it facing us? Or other people in some cases.
- ~117 Peers (single sw), 14 sending us BPDUs.
- Weird: two peers sending BPDUs from L3 interfaces. [vendor specific]
- Note: BPDUs are CPU processed, not line card processed.

FF: bcast

- Always interesting to see what ends up here.
- Normally at the IX we should only see ARP who-has.

ARP: 0x0806/FF

Filtering	Count
Total number of arp requests (who-has) in a 24 (HR)	3.3 - 3.4 million
Total number of arp requests to disconnected peers	629k – 786k
Valid requests	2.6 – 2.7 million
Invalid requests (percentage)	19-22%
Arp requests that are not part of 206.108.34.x/23?	3

MOP: 0x6002

- Maintenance Operation Protocol
  - Remote loading of software
  - Loopback testing
- MOP RC: layer 2 frame
- Enabled by default on many routers/ethernet interfaces
- Removing ***transport input mop*** from the VTY lines will **not** disable the MOP RC functionality
- Bonus: Linux tools latd/moprc

FF: bcast DNS

```
root@snoopdog:/data/captures/snoop_ispsummit# tcpdump -nr listenAll_20131101.pcap 'ether
dst ff:ff:ff:ff:ff:ff and not arp'
IP 206.108.34.YYY.53520 > 255.255.255.255.53: 308+ A? www.google.com
IP 206.108.34.YYY.57661 > 255.255.255.255.53: 309+ PTR? 33.158.104.ppp.in-addr.arpa.
IP 206.108.34.YYY.49531 > 255.255.255.255.53: 310+ PTR? 137.40.54.qqq.in-addr.arpa.
IP 206.108.34.YYY.51244 > 255.255.255.255.53: 311+ PTR? 165.40.54.rrr.in-addr.arpa.
IP 206.108.34.YYY.56892 > 255.255.255.255.53: 312+ PTR? 181.27.54.sss.in-addr.arpa.
IP 206.108.34.YYY.51672 > 255.255.255.255.53: 313+ PTR? 5.42.54.ttt.in-addr.arpa.
IP 206.108.34.YYY.55308 > 255.255.255.255.53: 314+ PTR? 22.29.54.uuu.in-addr.arpa.
IP 206.108.34.YYY.58881 > 255.255.255.255.53: 315+ PTR? 10.5.54.vvv.in-addr.arpa.
IP 206.108.34.YYY.54027 > 255.255.255.255.53: 316+ PTR? 130.103.104.www.in-addr.arpa.
IP 206.108.34.YYY.59046 > 255.255.255.255.53: 317+ PTR? 254.103.104.xxx.in-addr.arpa.
IP 206.108.34.YYY.52373 > 255.255.255.255.53: 318+ PTR? 120.254.85.yyy.in-addr.arpa.
IP 206.108.34.YYY.56727 > 255.255.255.255.53: 319+ PTR? 150.240.85.zzz.in-addr.arpa.
```

RWHO: ~0x0800/FF

PC routers

Some (harmless) fun with this peer

RWHO: ~0x0800

1. Download source code.
2. Recompile with changes.
3. Send ‘suggestions’ to a peer.

```
root@snoopdog:/tmp/netkit-rwho-0.17/rwhod# rwho -a
nistor   torix-says-turn-off-rwho:pts/1 Aug 26 23:35  1:36
nistor   torix-says-turn-off-rwho:pts/2 Aug 27 00:14
p?u?     tr!!l:pts/0                    Aug 18 22:35 99:59
p?u?     tr!!l:pts/2                    Aug 19 00:26 99:59
p?u?     tr!!l:pts/7                    Aug 19 00:29 99:59
ph?li?   tr!!l:pts/10                   Aug 22 13:13   :09
ph?li?   tr!!l:pts/11                   Aug 23 22:54 26:24
ph?li?   tr!!l:pts/14                   Aug 25 15:33 33:37
ph?li?   tr!!l:pts/9                    Aug 19 12:20 10:31
root@snoopdog:/tmp/netkit-rwho-0.17/rwhod#
```

NTP: It's what time?

TorIX: "We saw some issues around 20:09 EDT, could you quickly check your side, any errors you notice in the logs? Can you confirm the time?"

Peer: "I can't tell" or "yes we have errors, not sure on time"

```
33w6d: %LINEPROTO-SP-5-UPDOWN: Line protocol on
Interface FastEthernet1/36, changed state to down
33w6d: %LINK-3-UPDOWN: Interface FastEthernet1/36,
changed state to up
33w6d: %LINK-SP-3-UPDOWN: Interface FastEthernet1/36,
changed state to up
33w6d: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet1/36, changed state to up
```

NTP: It's what time?

router# **show clock detail**

*02:58:19.871 UTC Sun Mar 27 2011

Time source is hardware calendar

Hardware calendar, fabulous right? Nope!

LLDP: 0x88CC

```
14:09:05.562344 AABBCCDDEEFF > 01:80:c2:00:00:0e, ethertype LLDP (0x88cc), length 239
        Chassis ID TLV (1), length 7
          Subtype MAC address (4): 40:55:39:7d:0d:9e
        Port ID TLV (2), length 15
          Subtype Interface Name (5): Gi0/7/0/16.610
        Time to Live TLV (3), length 2: TTL 120s
        System Name TLV (5), length 27: TOROONXNPED01.PEER-NOT-NAMED.COM
        System Description TLV (6), length 87
          Cisco IOS XR Software, Version 4.1.1[Default]\0x0aCopyright (c) 2012 by Cisco
Systems, Inc.
        Port Description TLV (4), length 27: GigabitEthernet0/7/0/16.610
        System Capabilities TLV (7), length 4
          System  Capabilities [Router] (0x0010)
          Enabled Capabilities [Router] (0x0010)
        Management Address TLV (8), length 12
          Management Address length 5, AFI IPv4 (1): 142.xxx.yyy.2
          Interface Index Interface Numbering (2): 1540
        Management Address TLV (8), length 24
          Management Address length 17, AFI IPv6 (2): 260T:XXXX::YYYY:ZZZZ:1002
          Interface Index Interface Numbering (2): 1540
        End TLV (0), length 0
```

OSPF: 0x0800

```
23:53:19.566225 aabbccddeeff > 01:00:5e:00:00:05, ethertype IPv4 (0x0800), length 90:
(tos 0xc0, ttl 1, id 6406, offset 0, flags [none], proto OSPF (89), length 76)
    206.108.34.XXX > 224.0.0.5: OSPFv2, Hello, length 56 [len 44]
        Router-ID XXX.XXX.XXX.XXX, Backbone Area, Authentication Type: none (0)
        Options [External, LLS]
          Hello Timer 10s, Dead Timer 40s, Mask 255.255.254.0, Priority 1
          LLS: checksum: 0xfff6, length: 3
            Extended Options (1), length: 4
              Options: 0x00000001 [LSDB resync]
```

BGP: ~0x0800 / ~0x86DD

wrong prefix filters

wrong route-maps

ignoring your peers

BGP: wrong prefix list

Network	Next Hop	Path
30.0.0.0/22	206.108.34.xx	54321
31.0.0.0/19	206.108.34.xx	54321
66.0.0.0/24	206.108.34.xx	54321 3333
99.100.0.0/18	206.108.34.xx	54321 2222
101.255.0.0/20	206.108.34.xx	54321 **100 200** 32000
150.0.0.0/23	206.108.34.xx	54321 4444

BGP: wrong route-map

Network	Next Hop	Path
0.0.0.0	206.108.34.xx	54321
10.0.21.1/32	206.108.34.xx	54321
10.0.21.2/32	206.108.34.xx	54321
10.0.21.3/32	206.108.34.xx	54321
30.0.0.0/22	206.108.34.xx	54321
31.0.0.0/19	206.108.34.xx	54321
66.0.0.0/24	206.108.34.xx	54321 3333
99.100.0.0/18	206.108.34.xx	54321 2222

SNMP: In your network?

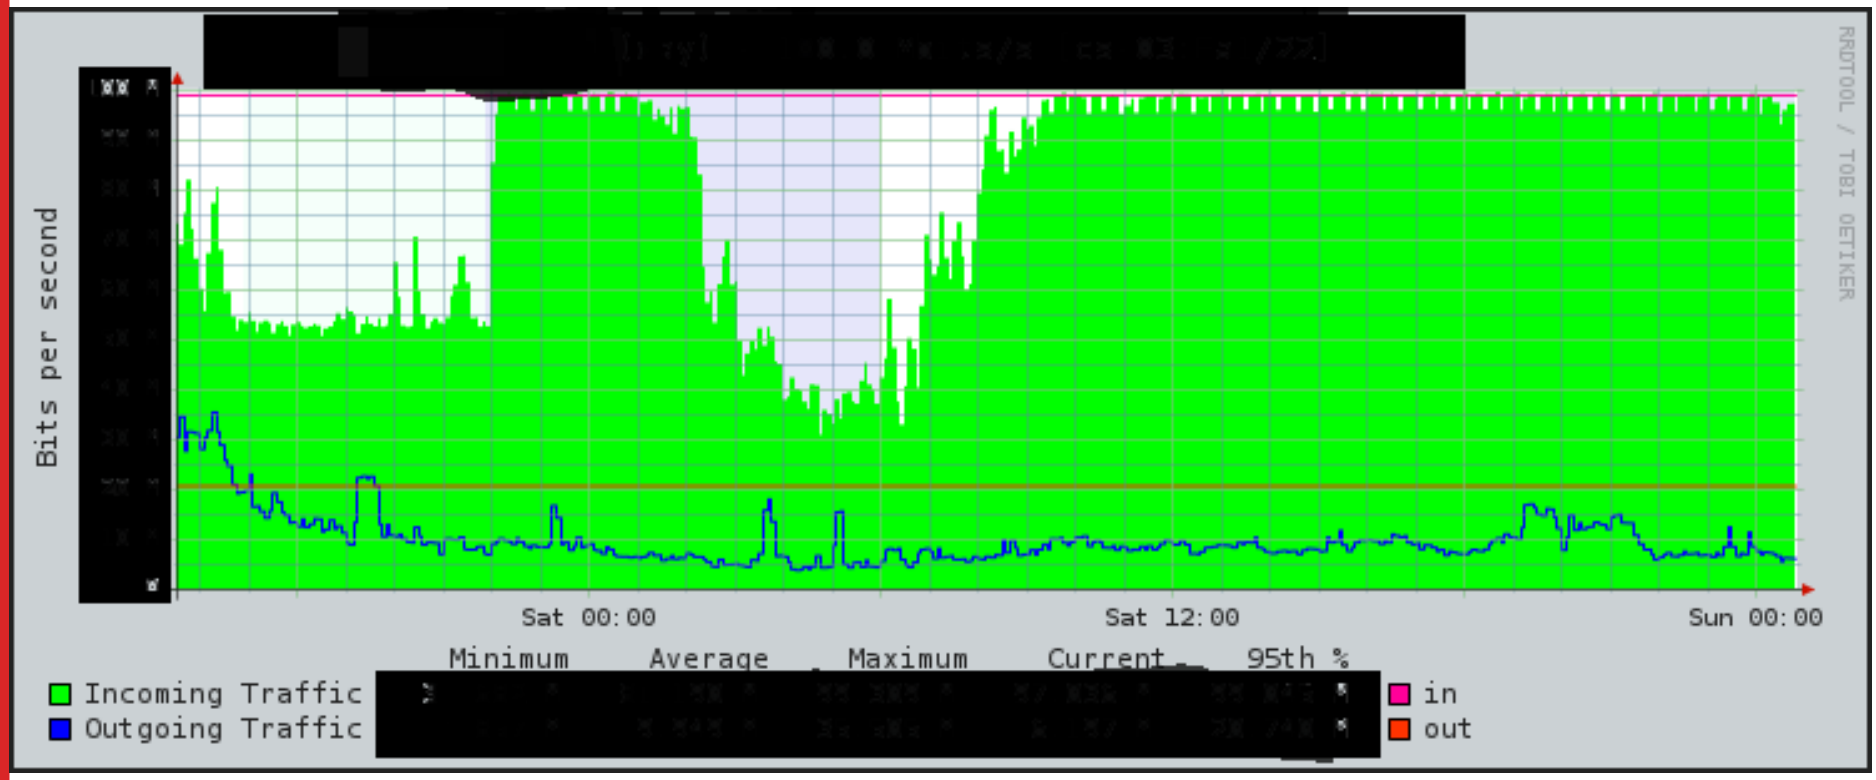

SNMP: errors in your network?

```
IF-MIB::ifInErrors.369098853 = 1364386973
IF-MIB::ifInErrors.436326400 = 1216281138
IF-MIB::ifInErrors.436830208 = 3756164
IF-MIB::ifInDiscards.436240384 = 511
IF-MIB::ifOutDiscards.437297152 = 6138
EtherLike-MIB::dot3StatsAlignmentErrors.436326400
= 1535
EtherLike-MIB::dot3StatsFCSErrors.436830208 =
3756162
```

Homework

Things you or your staff should do

.. and no excuses!

Homework

At a minimum

Check network for phy issues

Ensure all your devices are time sync'd

Turn off services you don't use

Are all of your neighbour sessions up?

(Does the peer even exist?)

Are you monitoring?

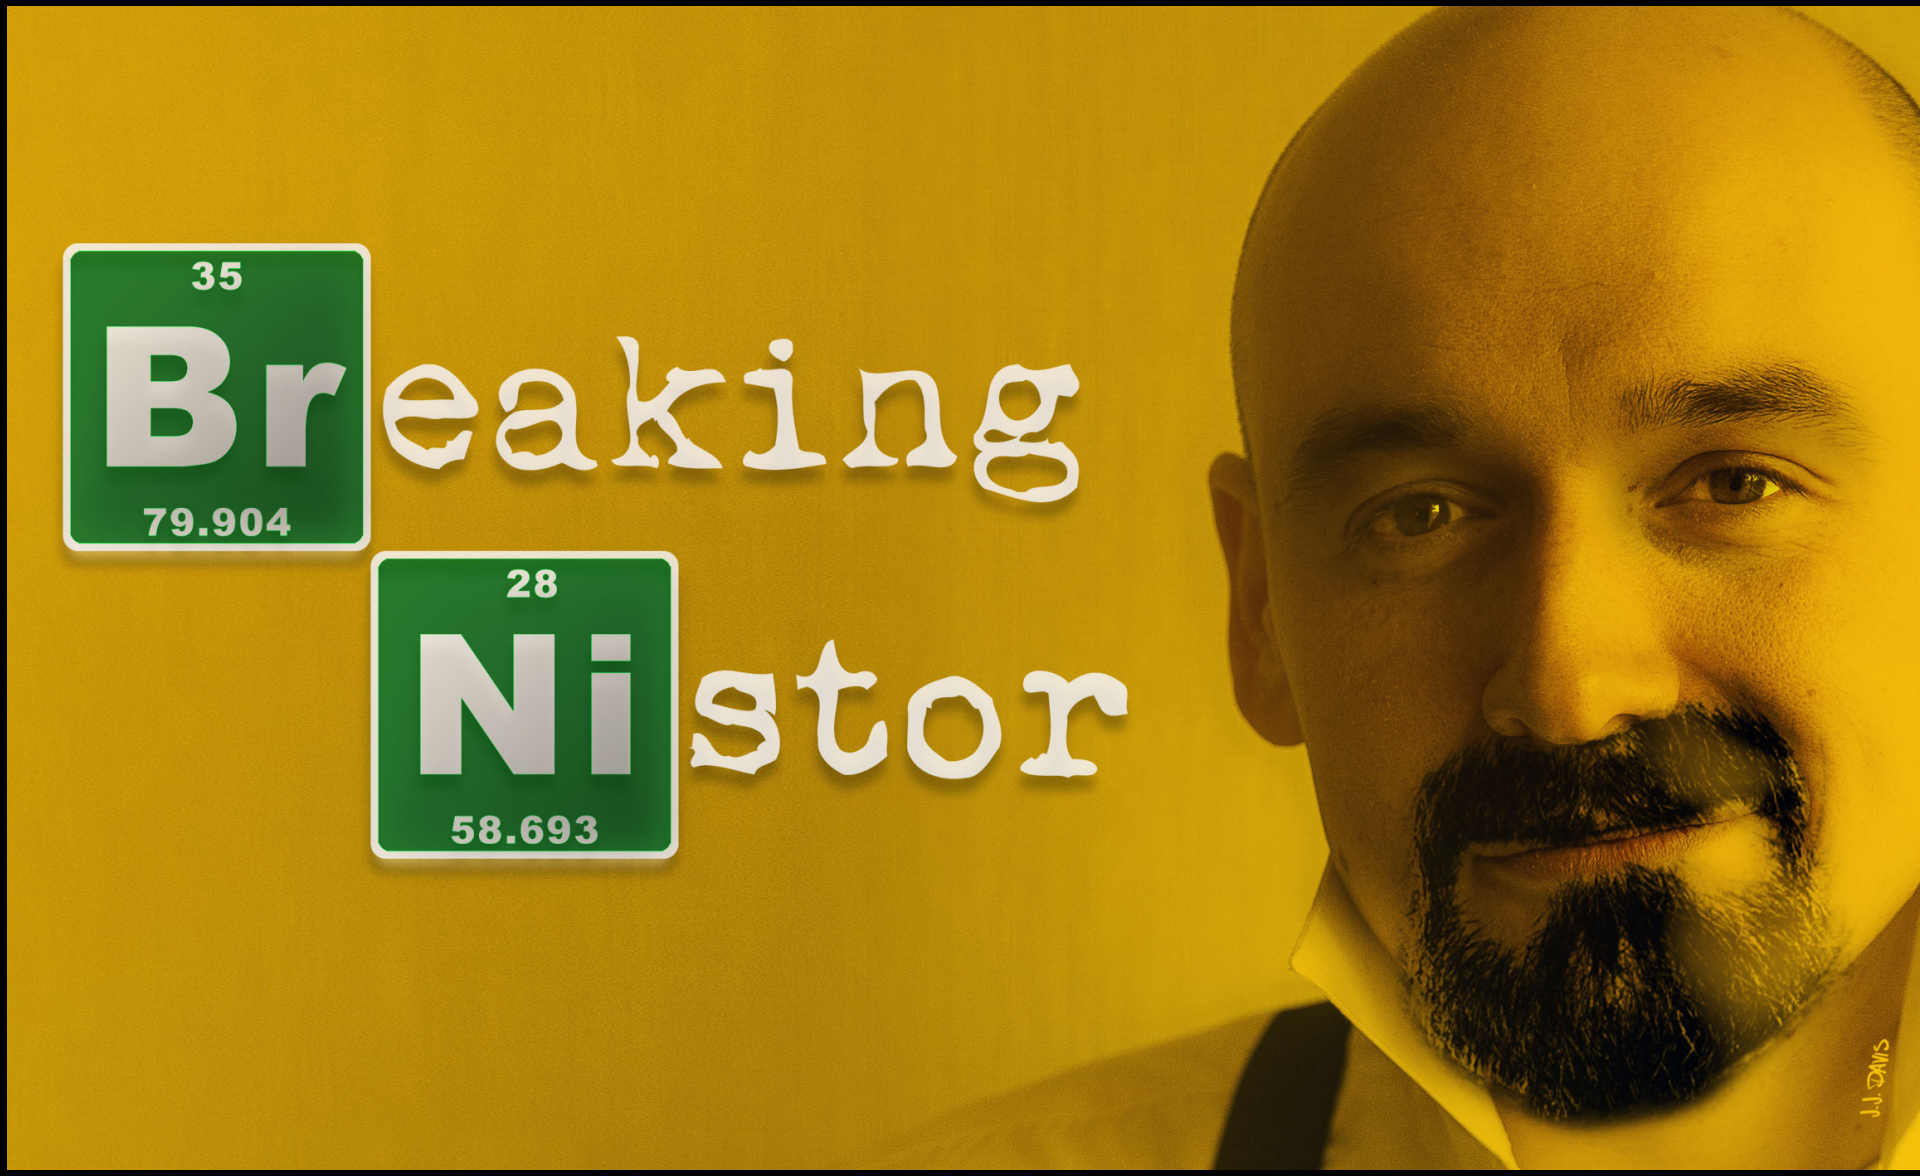
35
Br
79.904
eaking
28
Ni
58.693
stor