BERKELEY LAB
Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY
Office of Science

Network Tapping for Zeek

Michael Smitasin
Cyber Security Engineer
security@lbl.gov

October 14, 2022

Network Tapping?

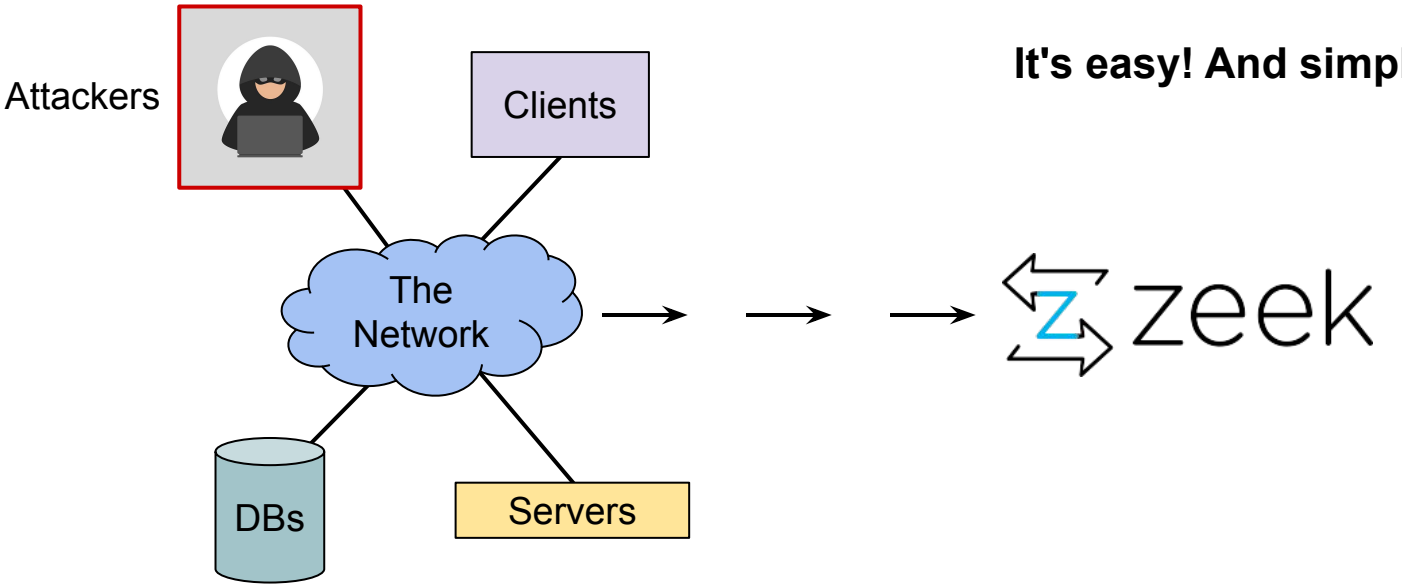

It's easy! And simple!

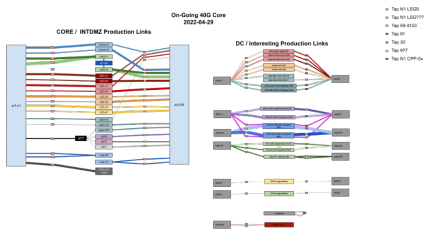

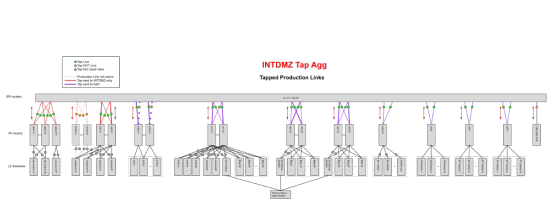

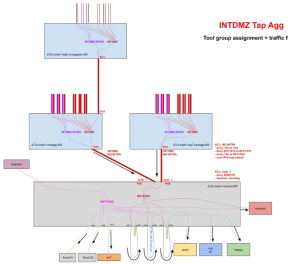

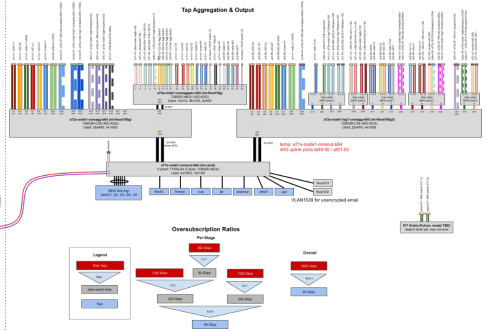

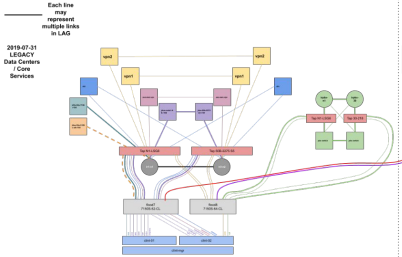

Except when it isn't

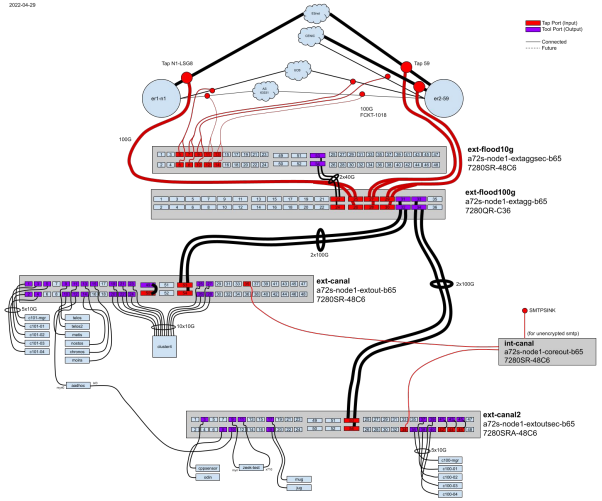

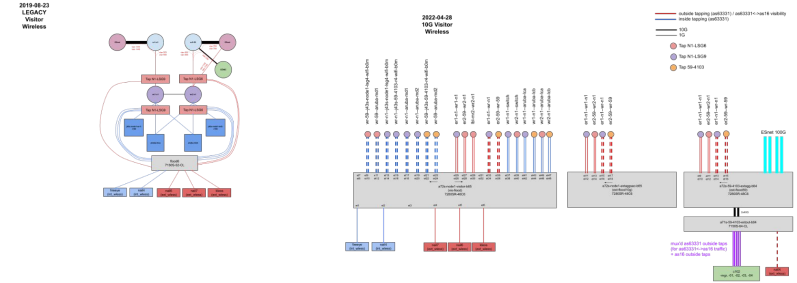

But really…

- You can do some cool things with taps and tap agg
- As long as you look out for the pitfalls

Mirror / Monitor / SPAN* Ports

- On-Device Packet Replication
- (+) Free?
- (+) Can filter at source
- (+) Non-disruptive add/change/remove
- (+) RSPAN/Lawful Intercept for remote capture
- (-) In-band / Resource contention?
- (-) Hardware limits
  - Ex: max 2 SPAN ports
- (-) Potential oversubscription
  - (1G TX, 1G RX = 2G tapped)

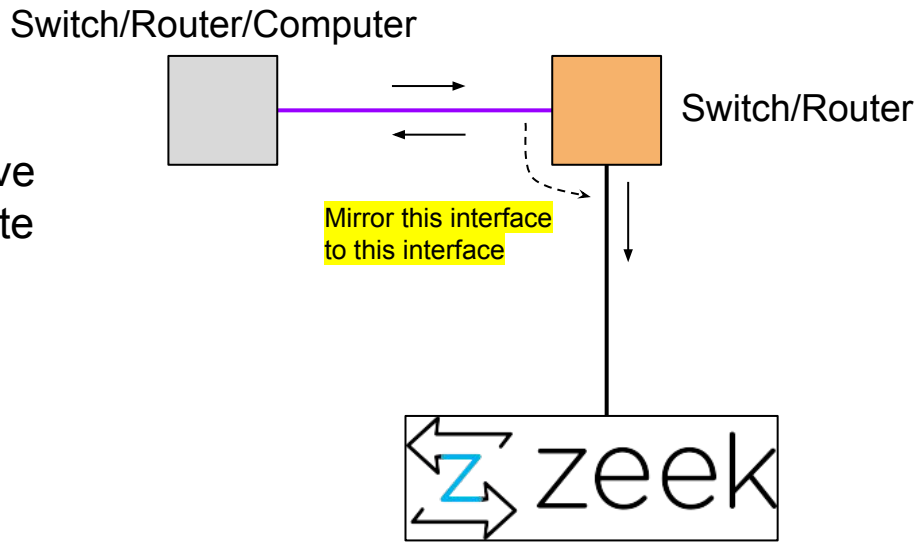

*Switch Port ANalyzer

Taps

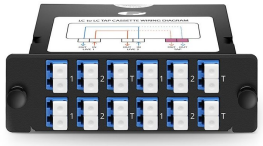

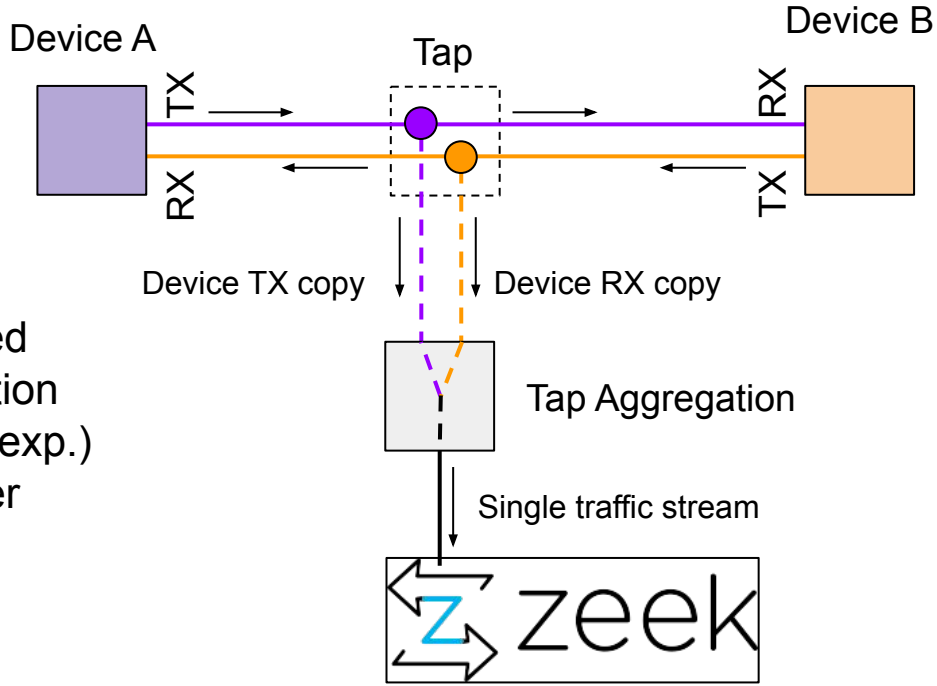

- (+) Out-of-Band
- (+) Fiber taps can be passive/unpowered
- (+) Fiber taps: all light, no oversubscription
- (+) Passive taps: Highly reliable (in our exp.)
- (-) Kinda expensive (cheaper than router ports?)
- (-) Disruptive add/change/remove

Fiber Pointers

- Connector types (LC vs MPO)

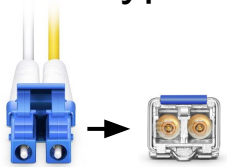

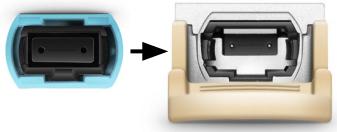

- Fiber types (Singlemode vs Multimode)

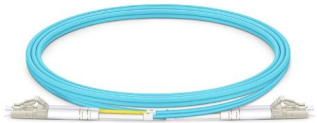

- Light Levels and Tap Split Ratios
- Good practices (bend radius & cleaning fiber)

Tap Agg Concepts (1)

- "Tap Agg Switch"
  AKA "Network Packet Broker"
- **Aggregate** taps to traffic streams
- **Filter** out traffic you don't want
- **Replicate** copies to different tools
- **Distribute** across cluster nodes

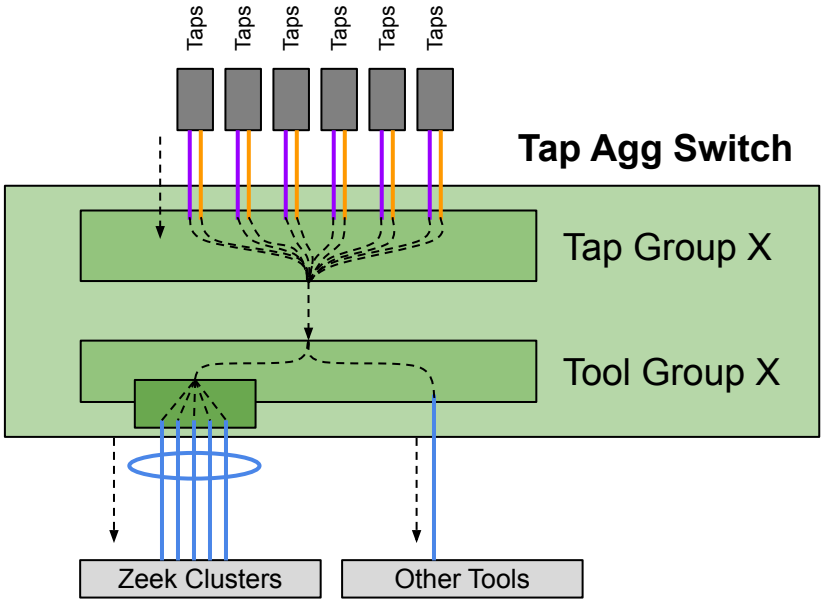

Tap Agg Concepts (2)

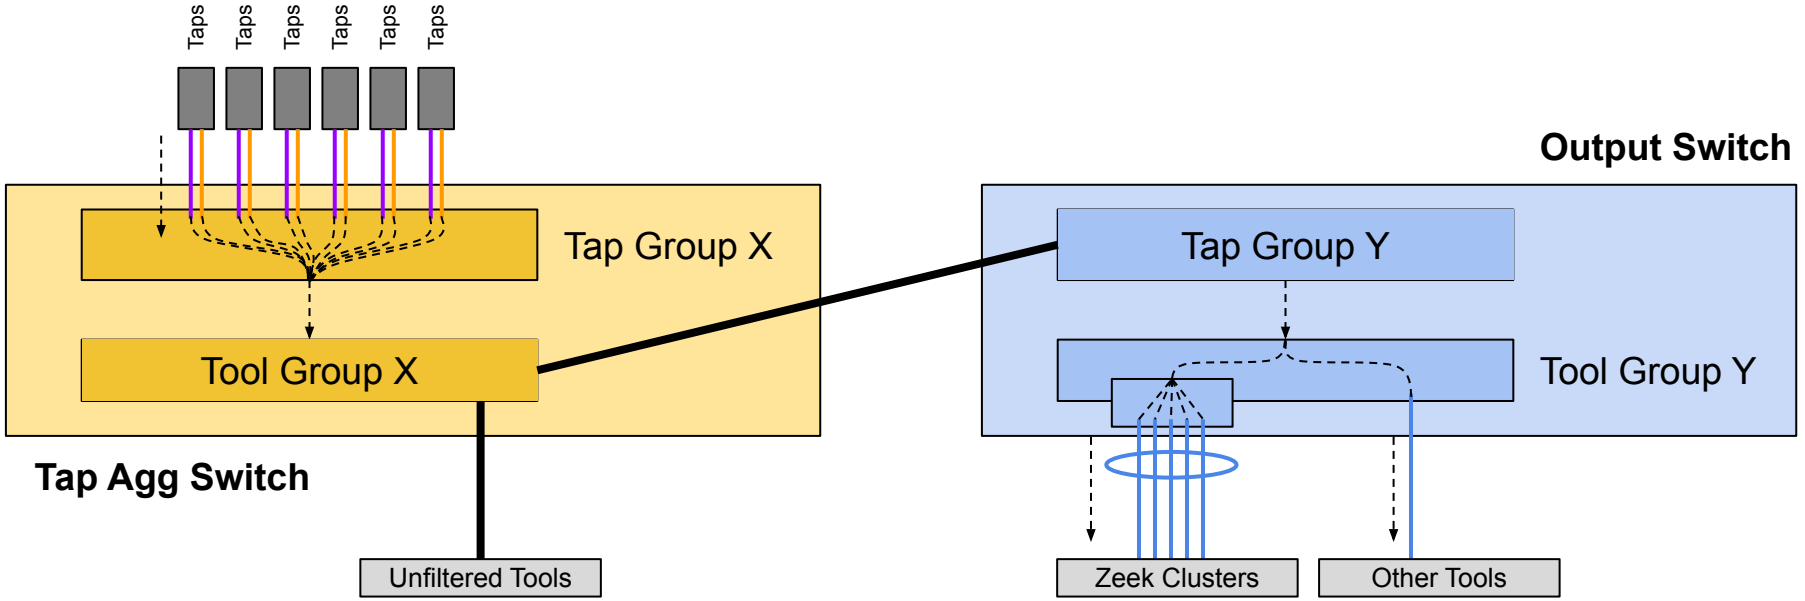

Distribute to a Zeek Cluster

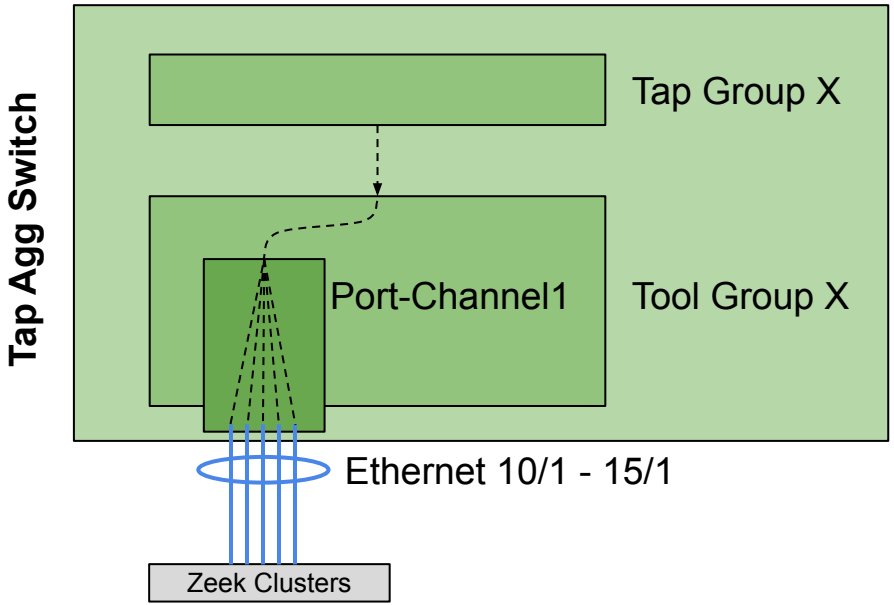


```
load-balance policies
  load-balance sand profile symmetric
    no fields mac
    fields ipv4 symmetric-ip
    fields ipv6 symmetric-ip
    fields l4 symmetric-ports
    no fields mpls
    fields symmetric-hash
    port-channel ip ip-tcp-udp-header

port-channel load-balance sand profile symmetric

interface Port-Channel1
  switchport mode tool
  switchport tool group set X

interface Ethernet 10/1 - 15/1
  channel-group 1 mode on
```

Distribute to Zeek Workers

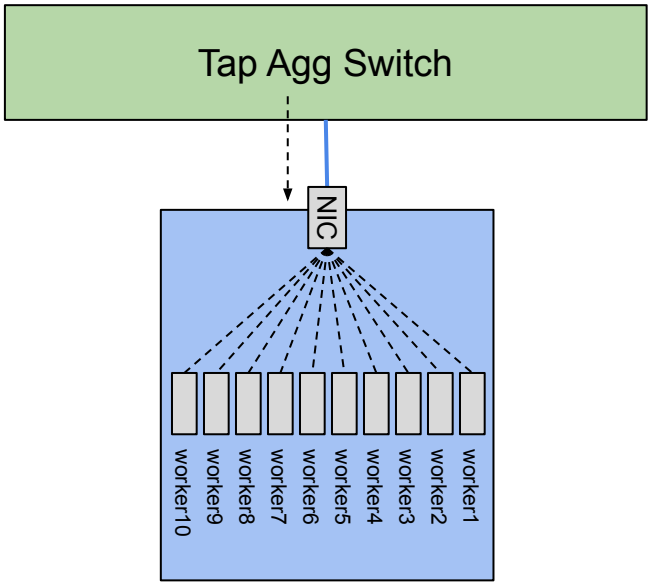

Zeek Node

```
# (zeekpath)/host/etc/node.cfg

# Myricom Sniffer Driver
lb_method=myricom
lb_procs=10
pin_cpus=3,5,7,9,11,13,15,17,19,21
env_vars=LD_LIBRARY_PATH=/usr/local/opt/snf/lib:/usr/local/
lib:$PATH,  SNF_DATARING_SIZE=0x80000000,
SNF_NUM_RINGS=10, SNF_FLAGS=0x1, SNF_APP_ID=1


# AF_Packet
lb_method=custom
lb_procs=10
pin_cpus=2,4,6,8,10,12,14,16,18,20
af_packet_fanout_id=23
af_packet_fanout_mode=AF_Packet::FANOUT_HASH
af_packet_buffer_size=128*1024*1024
```

Putting it all together

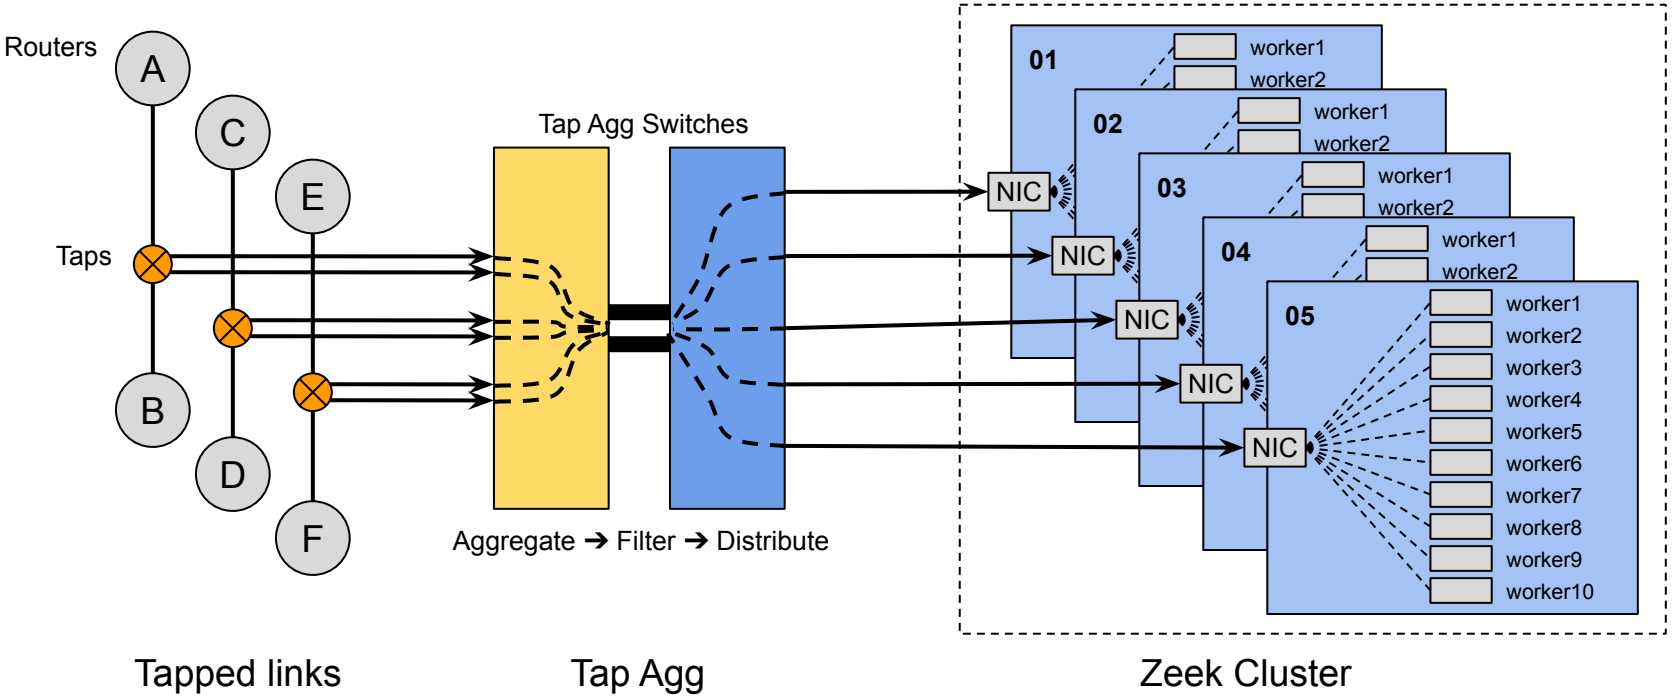

LBNL's Taps

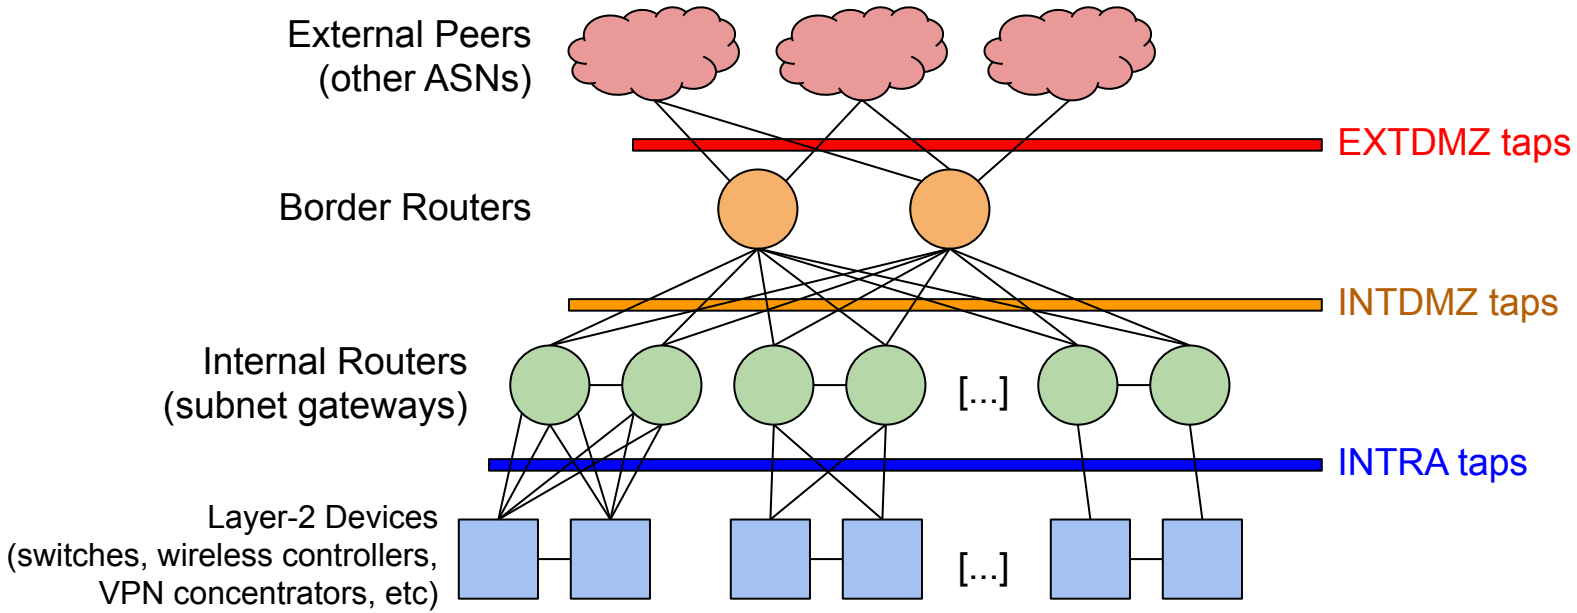
External Peers
(other ASNs)
EXTDMZ taps
Border Routers
INTDMZ taps
Internal Routers
(subnet gateways)
[...]
INTRA taps
Layer-2 Devices
(switches, wireless controllers,
VPN concentrators, etc)
[...]

LBNL's Tap Agg (EXT1)

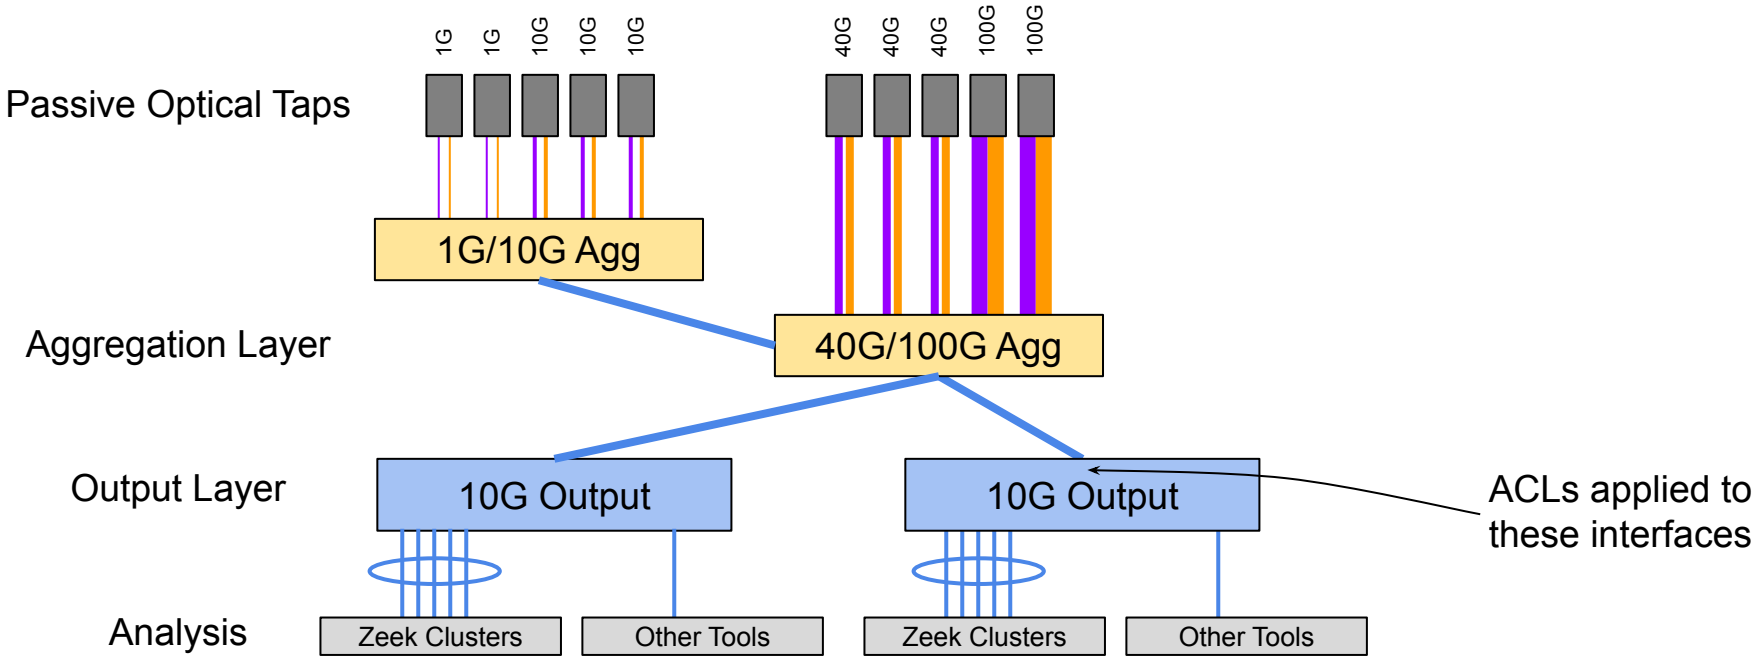

LBNL's Tap Agg (EXT2)

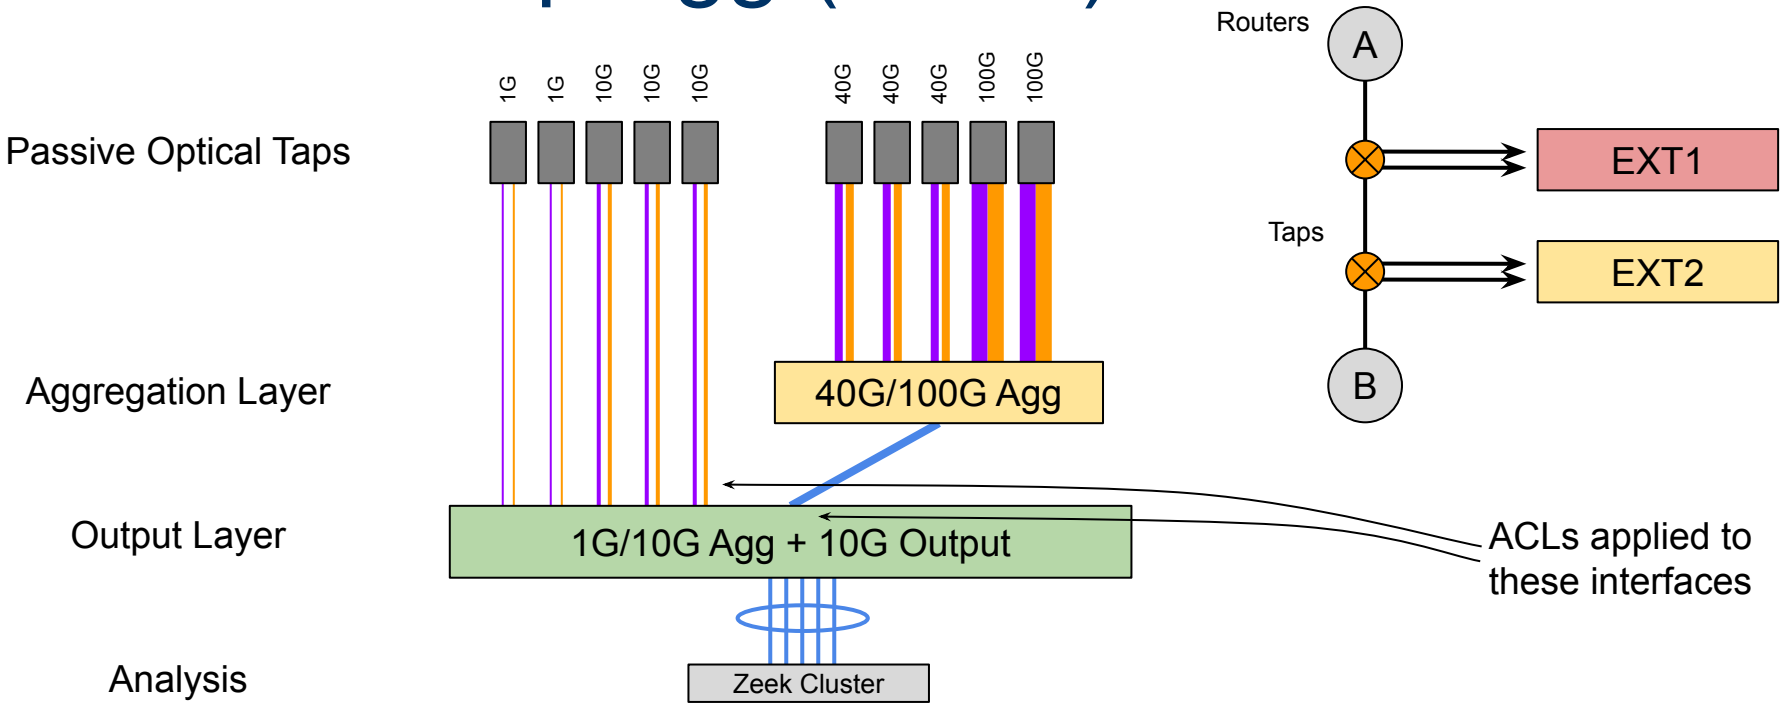

LBNL's Tap Agg (INTDMZ + INTRA)

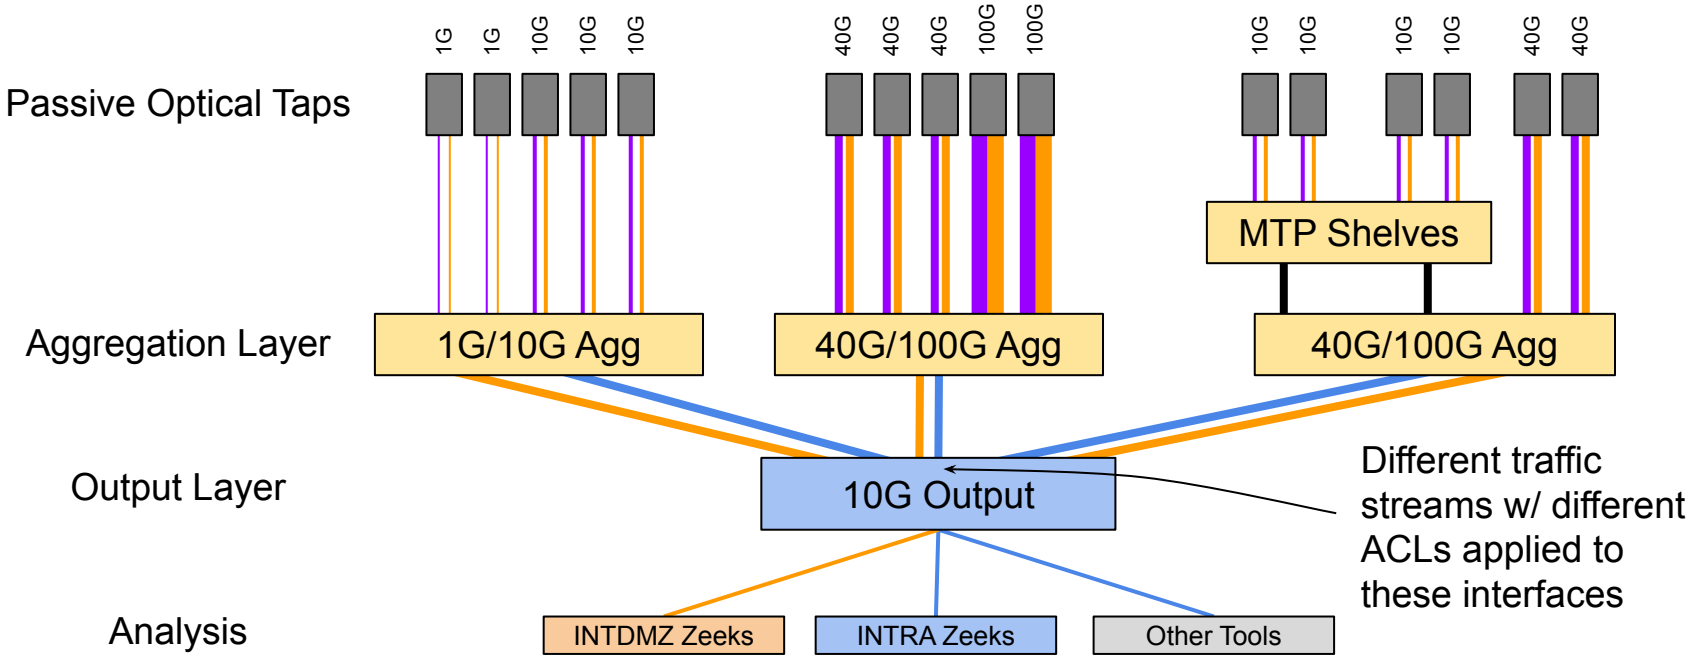

Static ACLing

- **Why: filter out specific traffic from being analyzed**
  - Protect low capacity tools

- First: Accept "Control Packets"
  - TCP SYN/FIN/RST + UDP + FRAG + GRE + ICMP
- Drop "fast-start" payloads
  - PerfSonar
  - xrootd
- Drop encrypted things when we get unencrypted too
  - SMTP (more later)

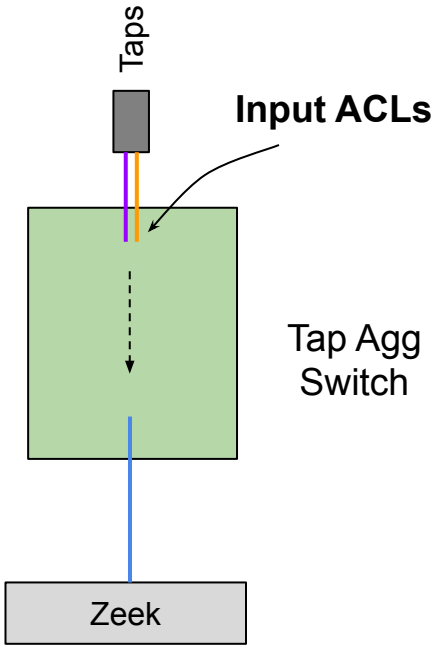

Static ACLing

TCP control packets + similar

```
ip access-list <ACLNAME>
   counters per-entry
   10 permit tcp any any syn
   20 permit tcp any any fin
   30 permit tcp any any rst
   40 permit tcp any any fragments
   50 permit udp any any
   60 permit gre any any
   70 permit icmp any any
   [...]
   100 deny ip host <perfsonar> any
   110 deny ip any host <perfsonar>
   [...]
```

PerfSonar Nodes

xRootd data transfer nodes

```
200 deny tcp any 131.243.135.0/26 range 1090 1100
210 deny tcp 131.243.135.0/26 range 1090 1100 any
220 deny tcp any range 1090 1100 131.243.135.0/26
230 deny tcp 131.243.135.0/26 any range 1090 1100
240 deny tcp any 131.243.135.0/26 range 10900 10910
250 deny tcp 131.243.135.0/26 range 10900 10910 any
260 deny tcp any range 10900 10910 131.243.135.0/26
270 deny tcp 131.243.135.0/26 any range 10900 10910
[...]
1000 deny tcp any host <SMTPSINK> eq smtp
1010 deny tcp host <SMTPSINK> eq smtp any
[...]
500001 permit ip any any
```

Encrypted SMTP

Dynamic ACLing

- **Can we do better than static ACLs?**
- Dynamically "shunt" big (elephant) flows' payloads
  - When you don't necessarily know what IPs/ports
- Still accept "control traffic" (like TCP SYN/FIN/RST)
- How: detect based on size threshold, add a 5-tuple ACL
  - Technically, can do other criteria
- conn-bulk.zeek -> dumbno.py -> API -> tap agg switch

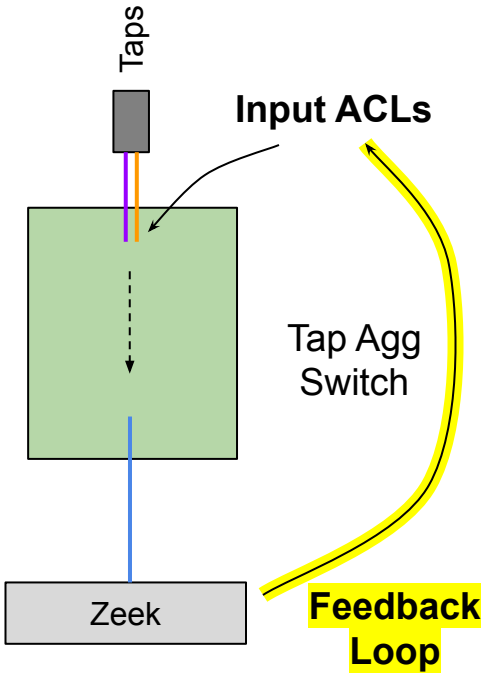

Dynamic ACLing :: ACL example

```
ip access-list bulk_1
   counters per-entry
   10 permit tcp any any fin
   20 permit tcp any any syn
   30 permit tcp any any rst
   40 permit tcp any any fragments
   50 permit udp any any
   60 permit gre any any
   70 permit icmp any any
   80 deny pim any any
[...]
   36075 deny tcp host 192.0.2.32 eq ssh host 203.0.113.5 eq 44144
   44051 deny tcp host 203.0.113.150 eq 62218 host 192.0.2.15 eq 50935
   44053 deny tcp host 203.0.113.150 eq 62220 host 192.0.2.15 eq 50935
   44057 deny tcp host 203.0.113.150 eq 62221 host 192.0.2.15 eq 50935
   44059 deny tcp host 203.0.113.150 eq 62222 host 192.0.2.15 eq 50114
   44623 deny tcp host 192.0.2.32 eq 53526 host 203.0.113.104 eq https
   45255 deny tcp host 192.0.2.116 eq 53042 host 203.0.113.188 eq https
[...]
   500001 permit ip any any
```

Accept TCP control packets + similar

Big Shunted Payloads

More Advanced: Selective Tapping

Missing Visibility

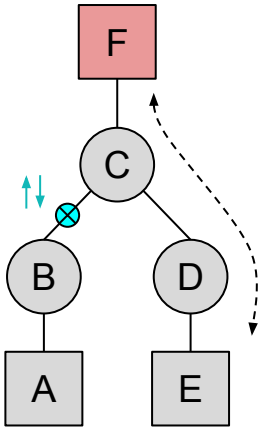

Too Many Copies

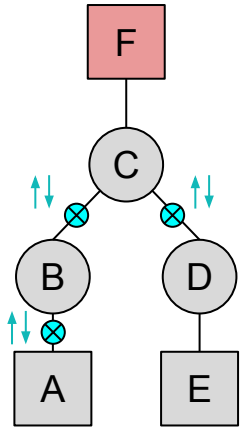

Just right

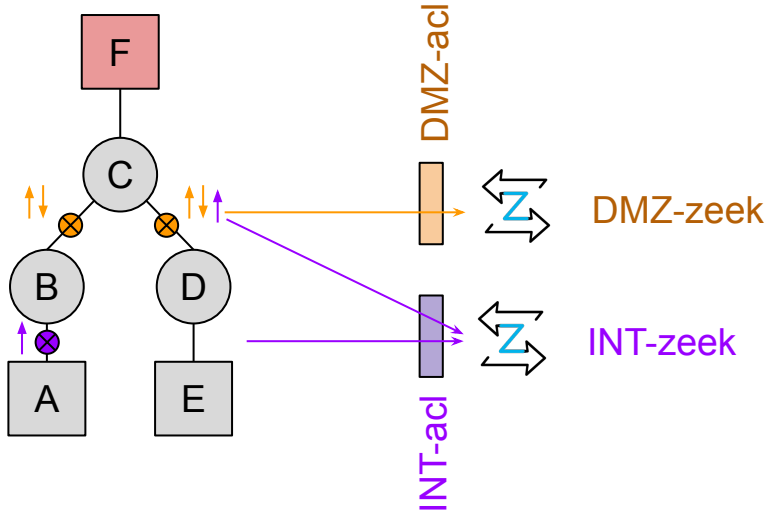

Tapping Email: Cloud+STARTTLS

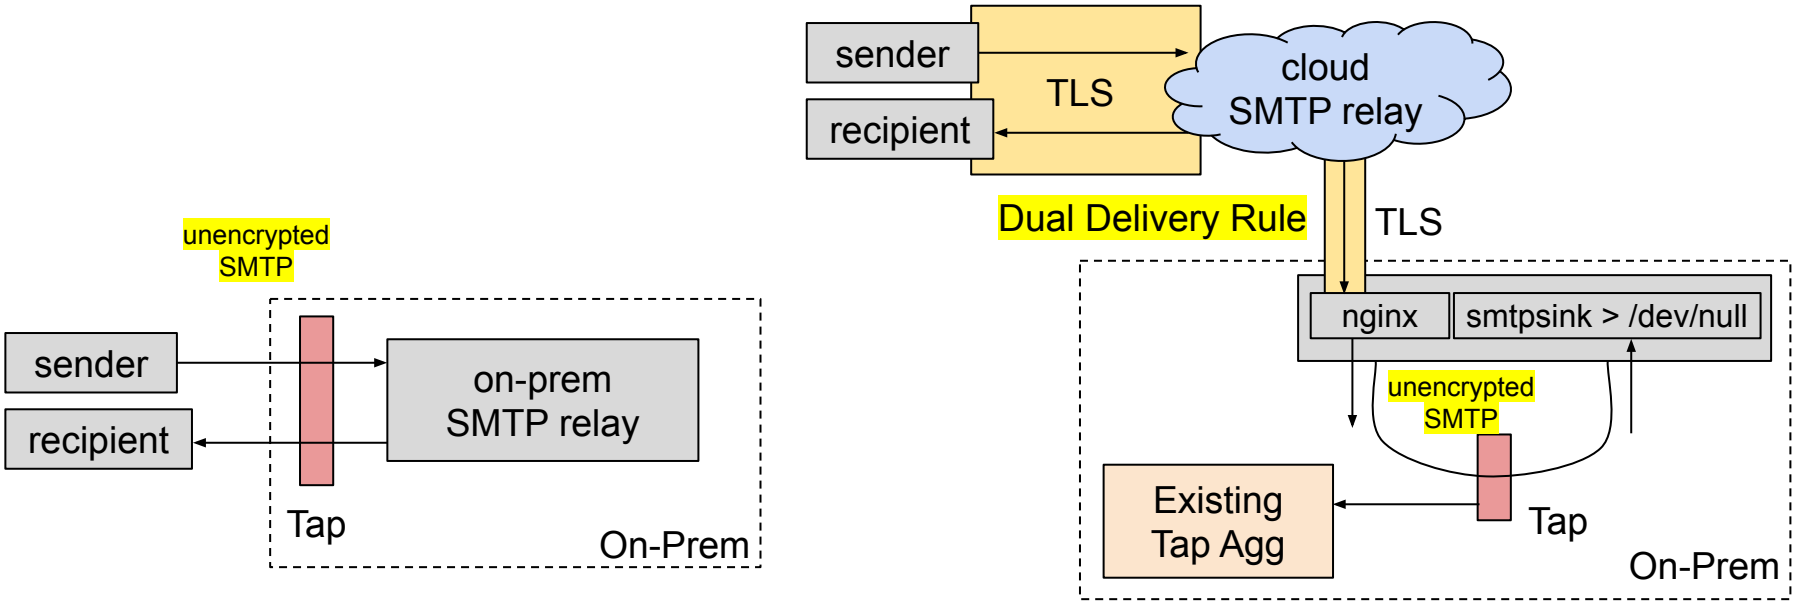

Appendix

- Calculating Light Budget
- Tap hardware
- Fiber Primer
- Fiber Connector Cleaners
- How to Install a Tap
- Checking Light Levels
- Tap Agg Hardware
- Hardware Example Install
- Minimum Tap Agg Config
- Dumbno Config / T-Shooting
- Zeek cluster hardware

BERKELEY LAB
Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY
Office of Science

Questions? Suggestions?

mnsmitasin@lbl.gov

security@lbl.gov

Appendix 1:

Calculating Light Budget

- Light split ratios: 50/50, 70/30, 80/20
  - Do you have enough light budget?

-0.5 = connector loss

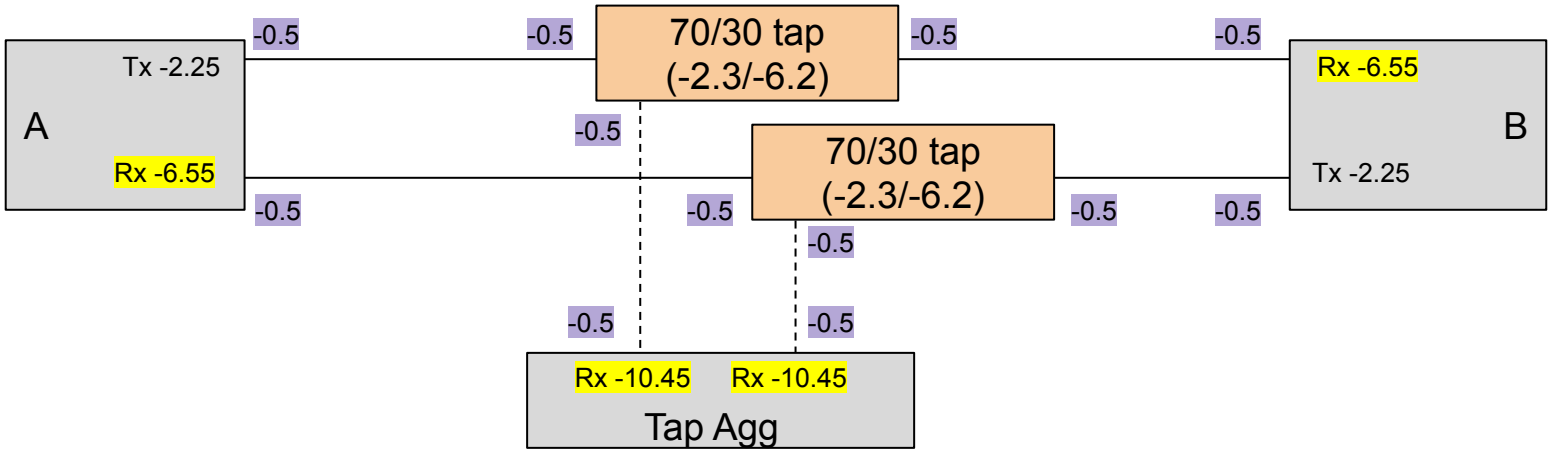

Appendix 2:

Calculating Light Budget

- Multiple taps for multiple locations

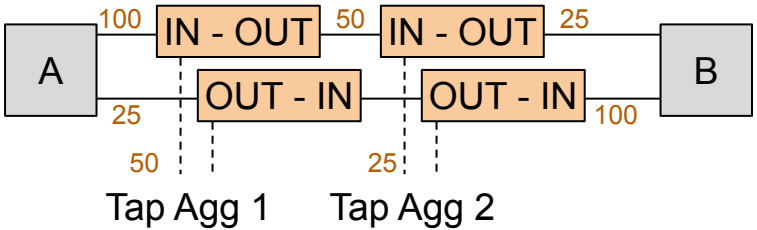

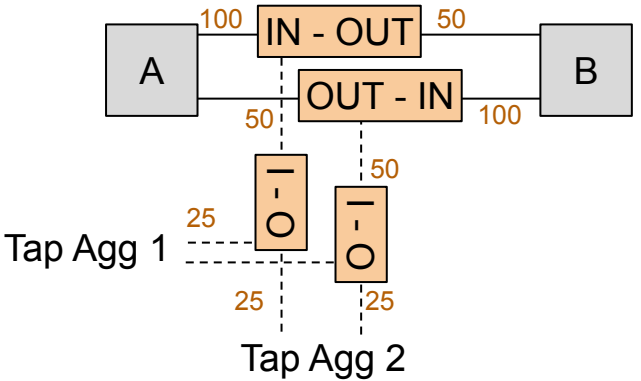

Appendix 3:

Tap Hardware

(1)

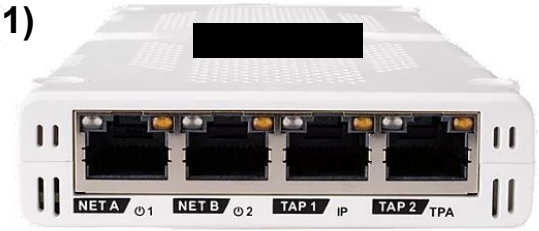

(2)

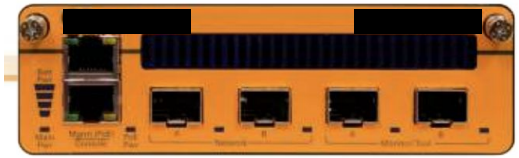

- Different Flavors of Taps
  1. Copper Taps
  2. Active Optical Taps
  3. Passive Optical Taps
  4. Fiber Patch Tap Cables

(3)

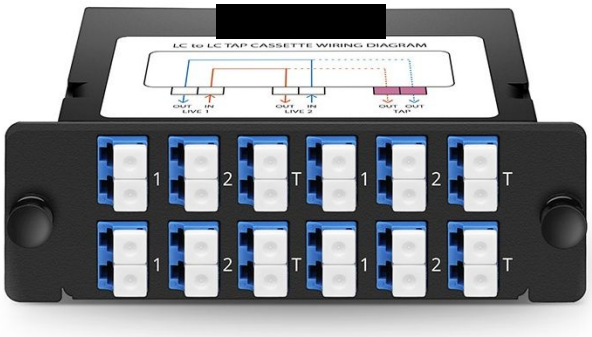

(4)

Appendix 4:

Fiber Primer

Common Fiber connectors

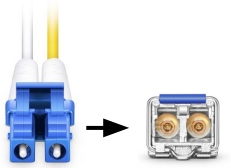

LC

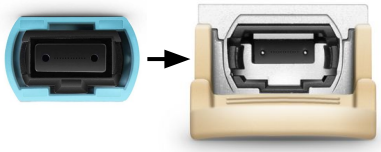

MPO (aka MTP®)

UPC vs APC
Don't mix these!

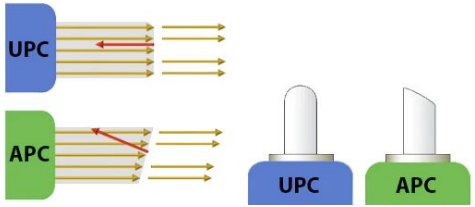

Common Fiber cables

OS2 (Optical Singlemode)
Long distance, any speed

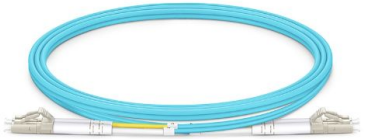

OM4 (Optical Multimode)
LC/LC connectors
Short dist., lower speeds

OM4 (Optical Multimode)
MPO-MPO connectors
(Polarity Type B)
Short dist., higher speeds

Appendix 5:

Fiber Connector Cleaners

Dirty connectors can cause link issues!

Appendix 6:

How to install a tap

1. Check with policy / legal counsel
2. Identify which specific link(s) you want to tap
3. Note the link type: copper/fiber, Singlemode/Multimode, connector type, speed (1G/10G/40G/100G)
4. Fiber: Check light levels, select appropriate ratio (80/20, 70/30, 50/50)
5. Plan what will plug-in where
6. Schedule a maintenance window (the link will go down)
7. Disconnect, clean connectors, add new cable, add tap
8. Confirm link comes up, check light levels after
9. Plumb the output to your tap agg or Zeek

Appendix 7:

Checking Light Levels

- Thresholds: device output below, or check the optical modules specs/data sheet, something like "Receiver Sensitivity" or "Receive Power" max/min.

- Cisco C6800s

```
#show interfaces Te1/1 transceiver detail
```

  "Optical Receive Power (dBm)"

- Arista 7280s

```
#show int et25/1 transceiver detail
```

  "Rx Power (dBm)"

- Juniper MX/EX

```
> show interfaces diagnostics optics et-1/0/2
```

  "Laser receiver power"

Appendix 8:

Tap Agg Hardware (1)

- Tap Agg Switches
  - Agg Layer:
    - 48x1G/10G + 6x40G/100G: [redacted]
    - 24x40G + 12x40G/100G: [redacted]
  - Output Layer:
    - 48x1G/10G + 6x40G/100G: [redacted]
  - "Tap Agg Mode" Licenses

- Zeek Node NICs
  - [redacted] 2x10G SFP+ w/ Sniff License (10G-PCIE2-8C2-2S+SNF3)
  - [redacted] 2x10G SFP+ w/ AF_Packet

Appendix 9:

Tap Agg Hardware (2)

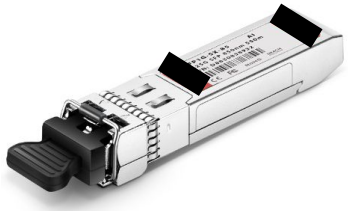

- Optical Modules
  - SFP = 1Gbps (most commonly, technically other s
  - SFP+ = 10Gbps
  - SFP28 = 25Gbps
  - QSFP+ = 40Gbps
  - QSFP28 = 100Gbps
  - QSFP-DD = 400Gbps
- Cables
  - DAC = Direct Attached Copper
  - AOC = Active Optical Cable

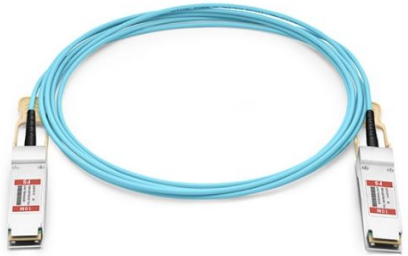

Appendix 10:

Hardware Example Install

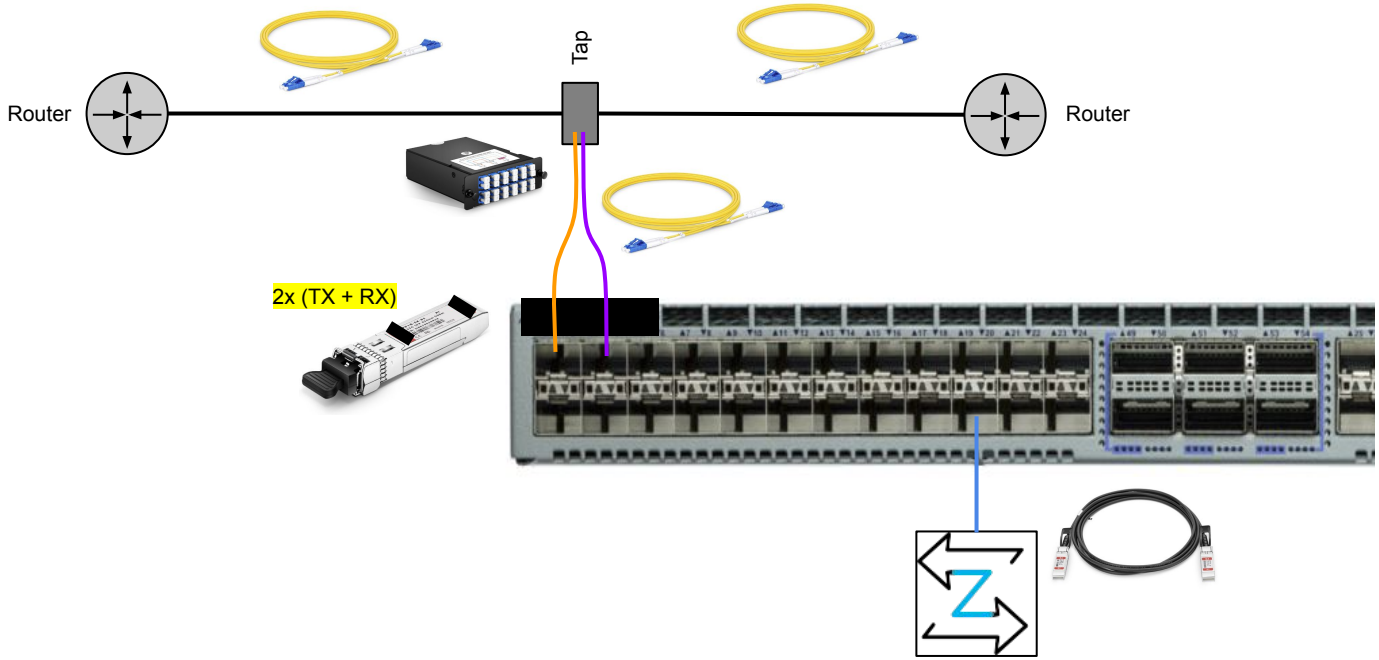

Appendix 11:

Minimum [redacted] Tap Agg Config

(No ACLs, no port channels to clusters)

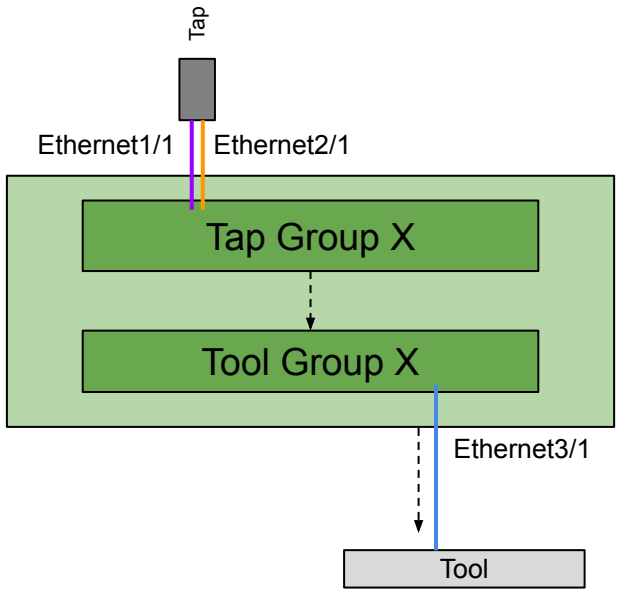


```
tap aggregation
  mode exclusive

interface Ethernet1/1
  description "TX Tap Input"
  switchport mode tap
  switchport tap default group X

interface Ethernet2/1
  description "RX Tap Input"
  switchport mode tap
  switchport tap default group X

interface Ethernet3/1
  description "Output to Tool"
  switchport mode tool
  switchport tool group set X
```

Appendix 12:

Dynamic ACLing :: conn-bulk.zeek

```
export {
    const size_threshold = 134217728 &redef; #128 megabytes
```

Connection cut-off in bytes

```
if ((( c$orig$size > size_threshold || c$resp$size > size_threshold ) && c$orig$num_pkts > 100 && c$resp$num_pkts > 100))
        event Bulk::connection_detected(c);
        return -1sec;
 }
```

You could use other criteria here too:

- orig_pkts
- resp_pkts
- IPs
- ports/protocols

Appendix 13:

Dynamic ACLing :: dumbno.cfg

```
[switch]
ip = <Tap Agg mgmt IP>
user = <APIUSER>
password = <APIPASSWORD>

[ports]
Ethernet1 = <Dynamic ACL name applied to ingress Tap ports>

[egress_ports]
Ethernet2 = tool1
```

Input port(s) from taps

Output port(s) that goes to Zeek

Appendix 14:

Dynamic ACLing :: T-Shooting logs

- Zeek :: conn_bulk.log

```
1663570498.392966 Coqv5l3qjHNZjqN1ag 192.0.2.70 44470 203.0.113.63 443 tcp ssl
1.688105 625 445138831 SF F T 0 ShADdFafRR 14 1197 8 2687 -
worker-2-1 LK US
```

- /var/log/dumbno/

```
@400000006323aaee121ea624 INFO:dumbno:op=ADD seq=32905 rule='tcp host 192.0.2.70 eq 44470
host 203.0.113.63 eq 443'
@400000006323aaf4267ad28c INFO:dumbno:op=REMOVE acl=bulk_1 family=ip seq=32905 rule='tcp
host 192.0.2.70 eq 44470 host 203.0.113.63 eq 443'
```

- /var/log/dumbno-stats/

```
@40000000632738e330d4639c INFO:dumbno_stats:mbps: in=3633 out=1852 filtered=1780
```

Appendix 15:

Zeek Cluster Hardware

- Zeek Cluster Nodes
  - (1x) Manager
    - [redacted] 2216RSJ2L-2T chassis
    - 2x [redacted], 20x cores @ 2.10GHz
    - 512GB (16x32GB) DDR4 RAM
    - 2x1TB NVMe (OS - [redacted], 4x3.8TB SSD (Data - [redacted])
    - 1x [redacted] 10G NIC
  - (4-5x) Worker Nodes
    - [redacted] 2216RSJ2L-2T chassis
    - 2x [redacted], 20x cores @ 2.10GHz
    - 256GB (8x32GB) DDR4 RAM
    - 2x1TB NVMe (OS - [redacted])
    - 1x [redacted] 10G NIC