Untrusting the network

Aijay Adams
Jose Leitao
Production Network Engineers

facebook

as of June 2016

1.13 billion daily active users on average

1.57 billion mobile monthly active users

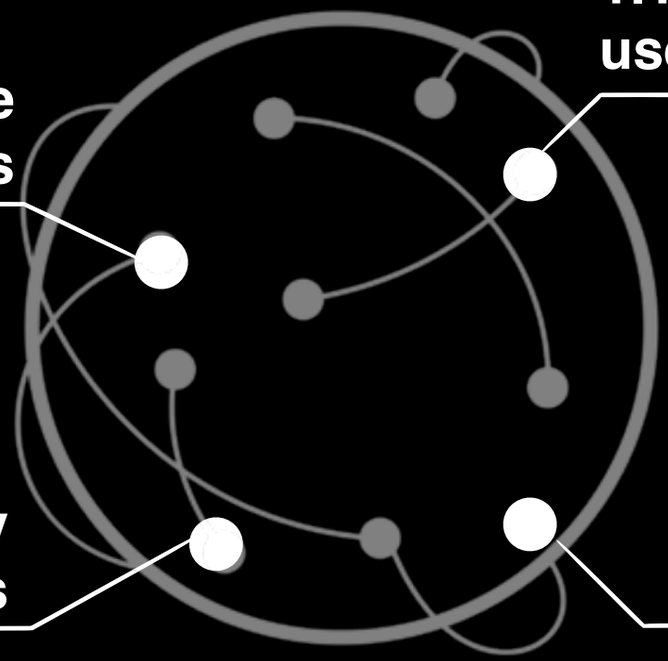

1.71 billion monthly active users

1.03 billion mobile daily active users on average

facebook

as of June 2016

1.13 billion daily active users on average

1.57 billion mobile monthly active users

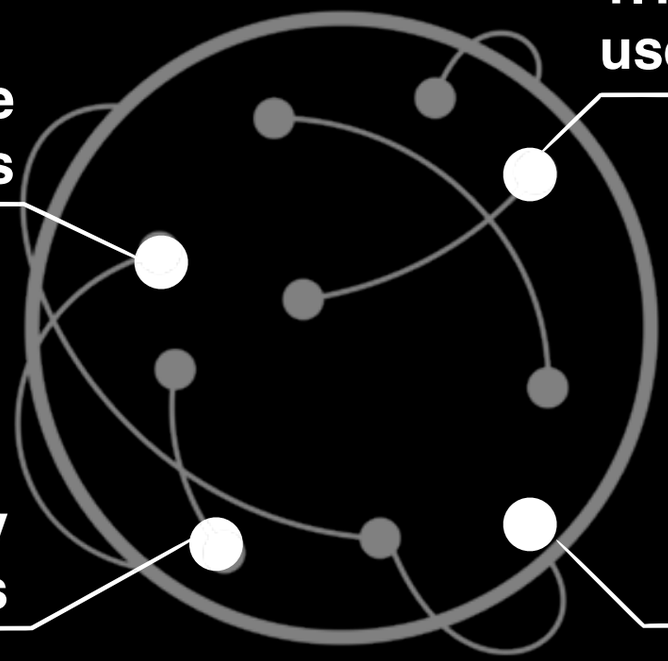

1.71 billion monthly active users

1.03 billion mobile daily active users on average

Approximately 84.5% of our daily active users are outside the US and Canada.

Data Center Network

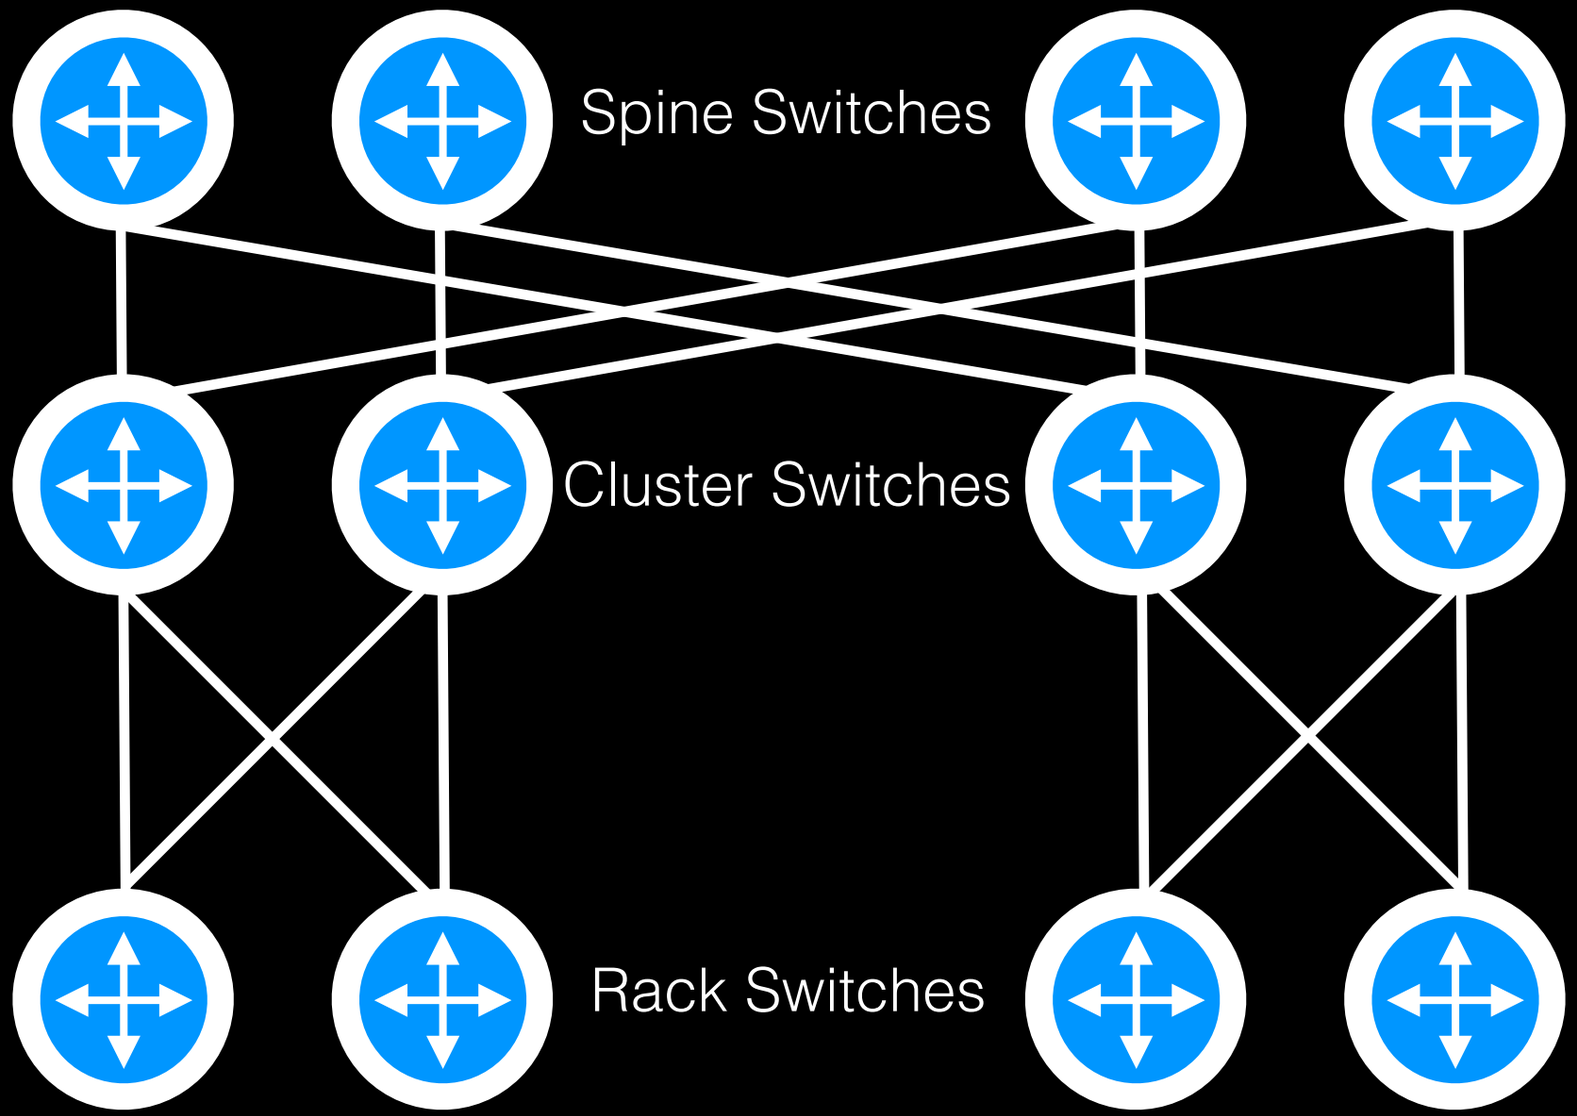

Wide ECMP, many paths!

Backbone Network

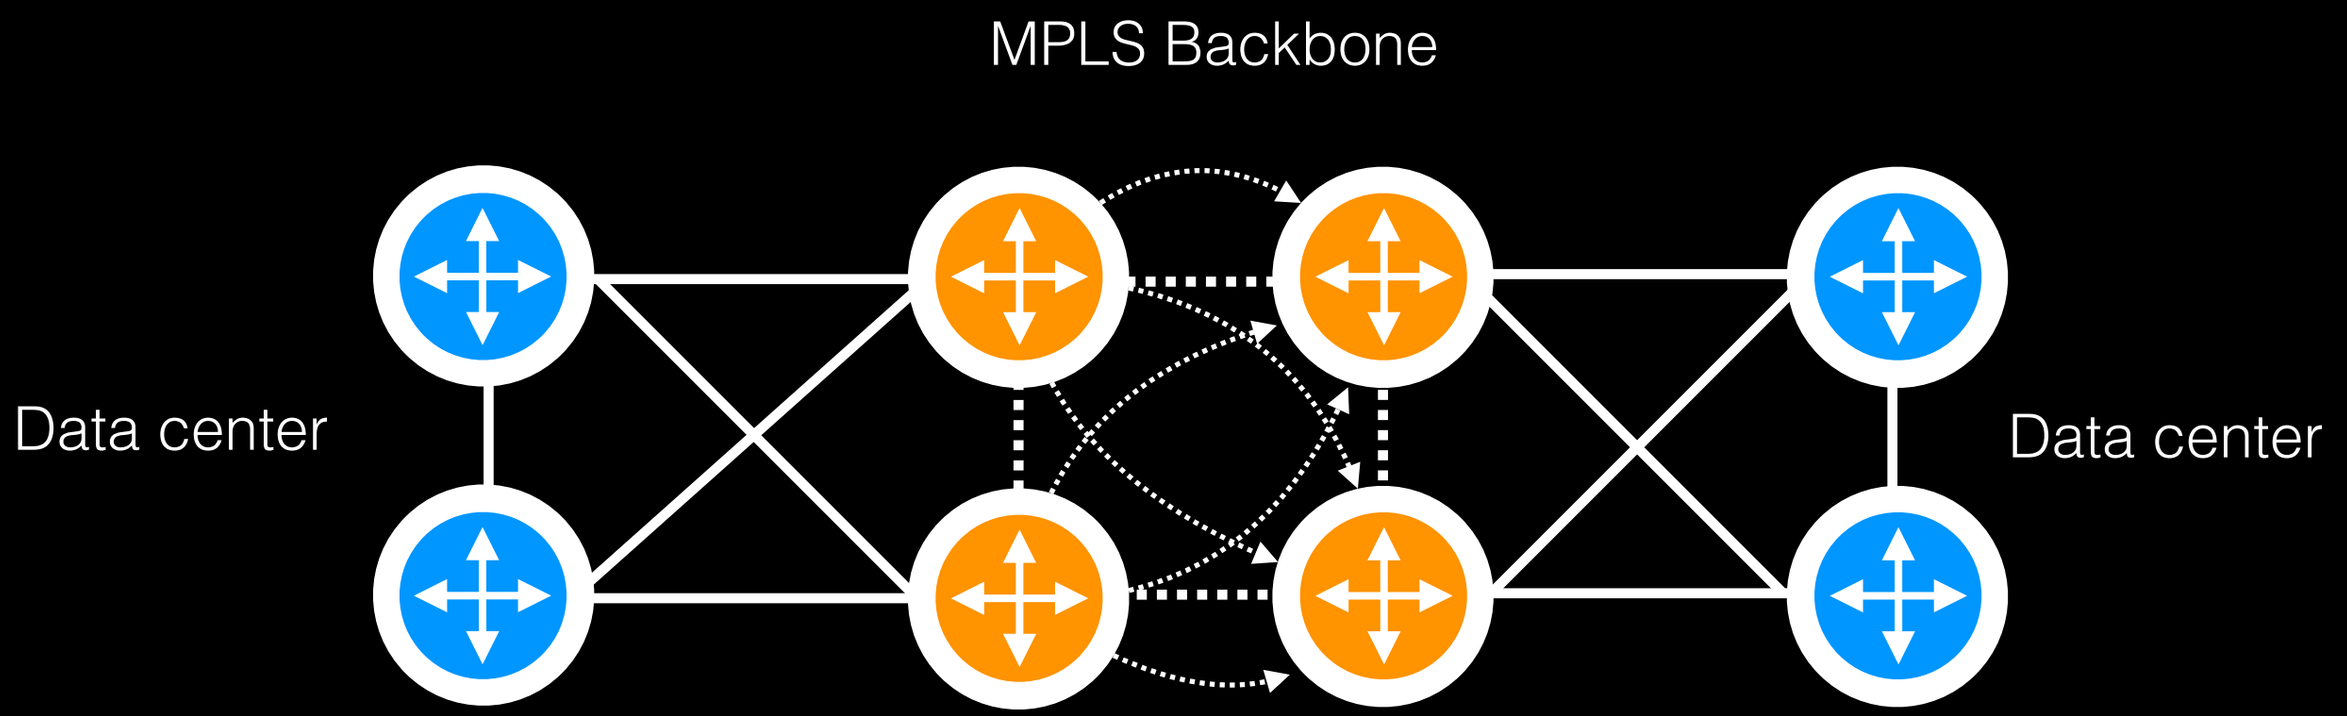

Auto Bandwidth

ECMP over MPLS Tunnels

Backbone

Region A

DC 1

Racks

Cluster 1

Cluster 2

Racks

DC 2

Cluster 1

Racks

POP B

Cluster 1

Racks

Monitoring the Network

?!??

Monitoring the Network

Counters and Logs

Monitoring the Network

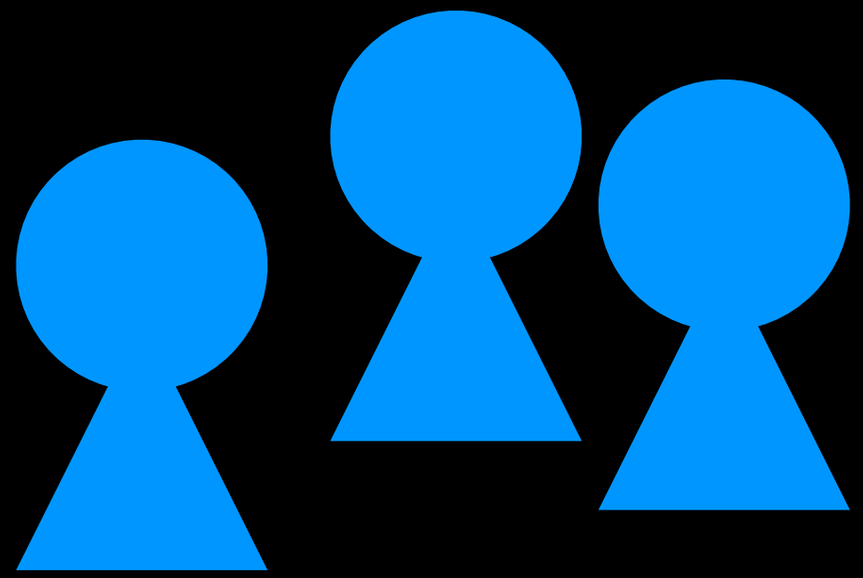

Coworkers

Monitoring the Network

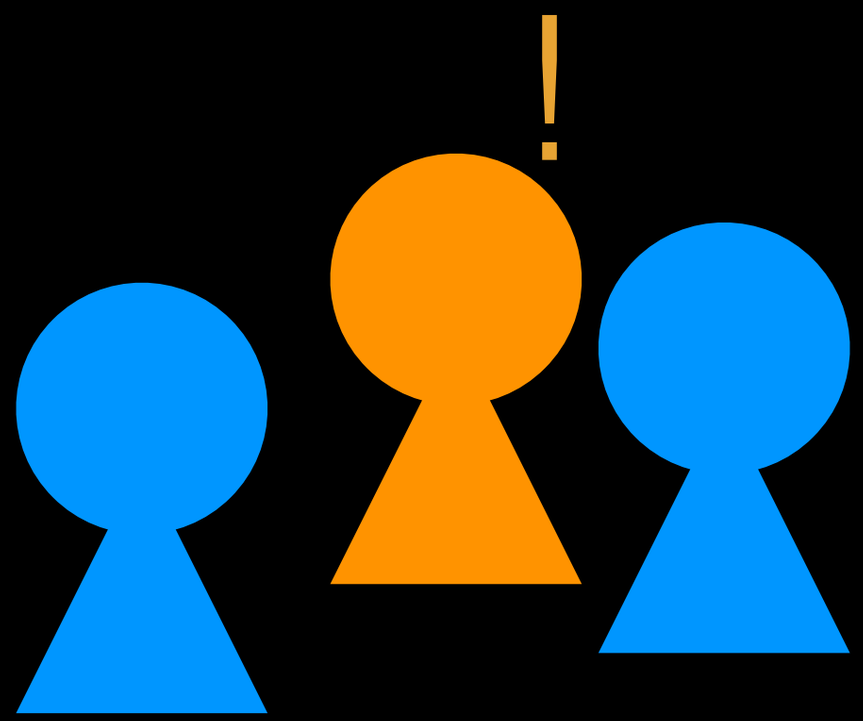

Coworkers

Investigate

#network_engineer

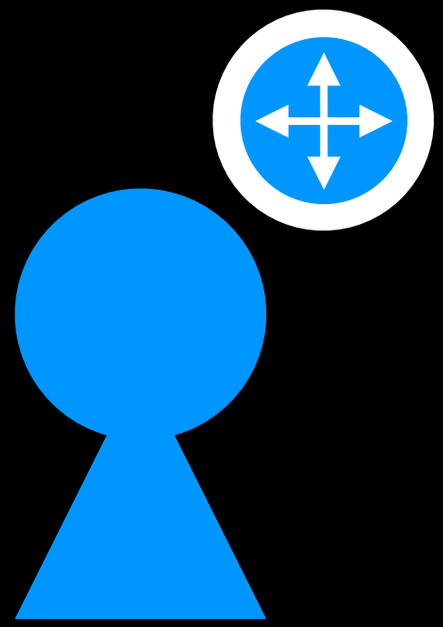

NetNORAD

The network fault detector

github.com/facebook/UdpPinger

Ping all the things!

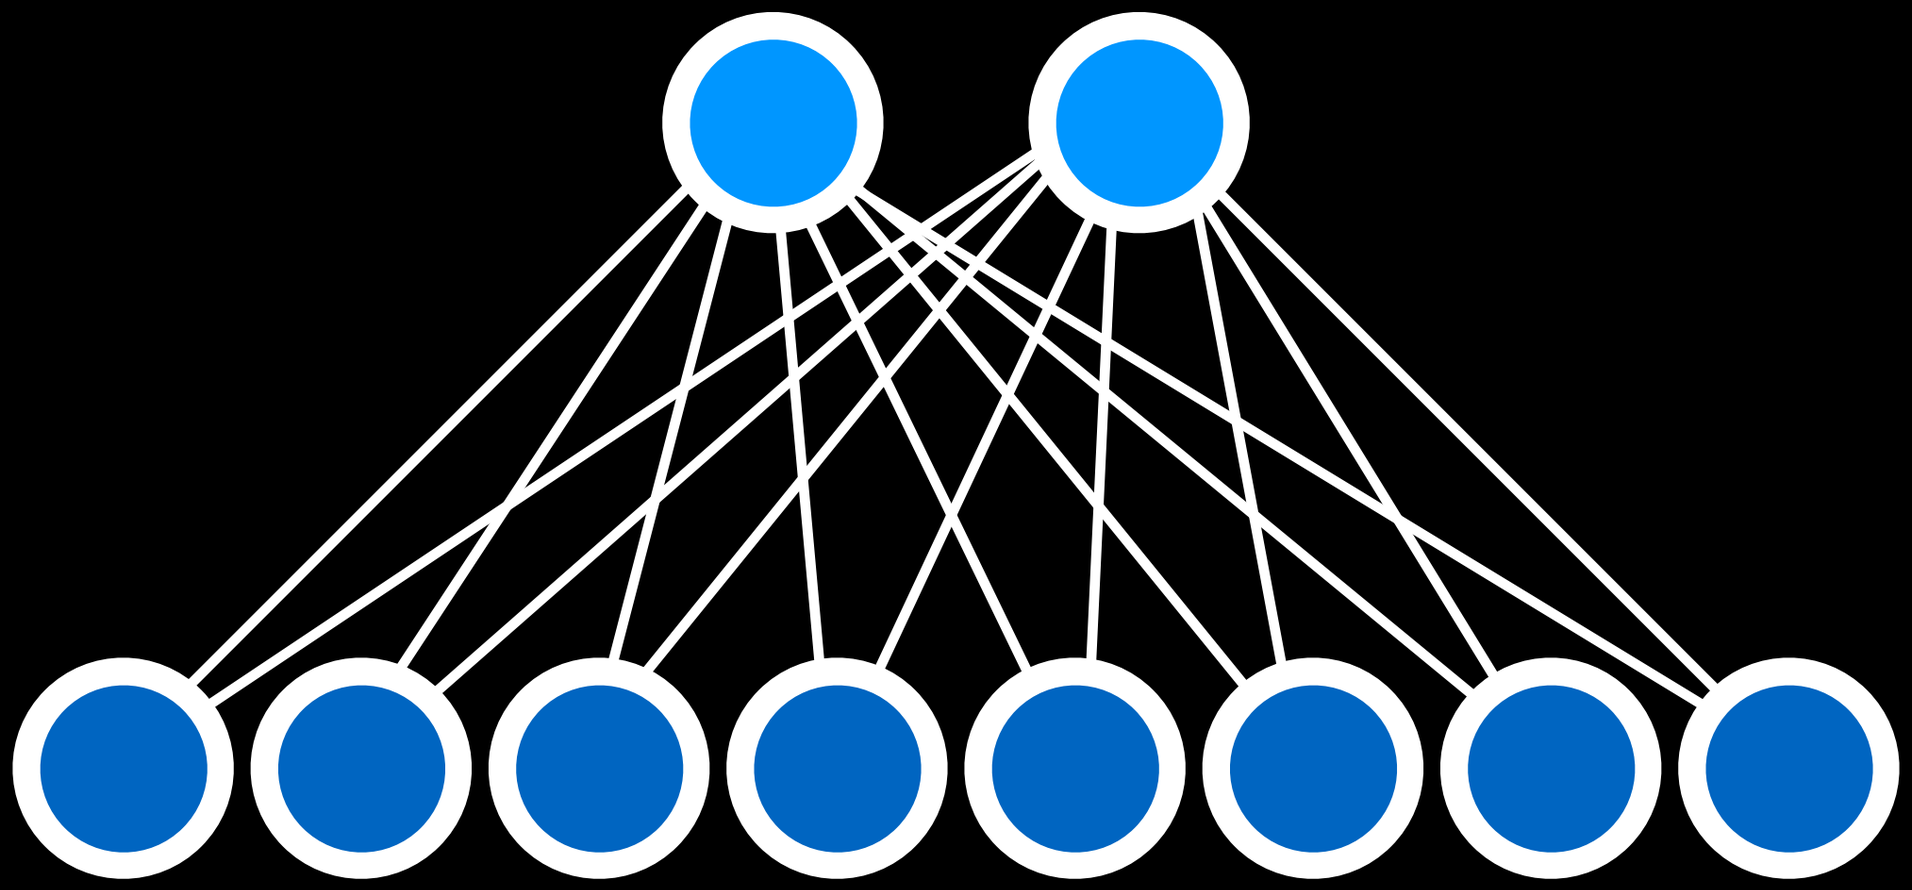

Run pingers on some machines
Run responders on all machines
Collect and analyse data

Evolution

- Run /bin/ping from a python agent
- Raw Sockets, Fast TCP Probes
- Raw Sockets, Fast ICMP Probes
- UDP Probes and Responder
- UDP Probes and Responder + Fast ICMP Probes

Ping Pong

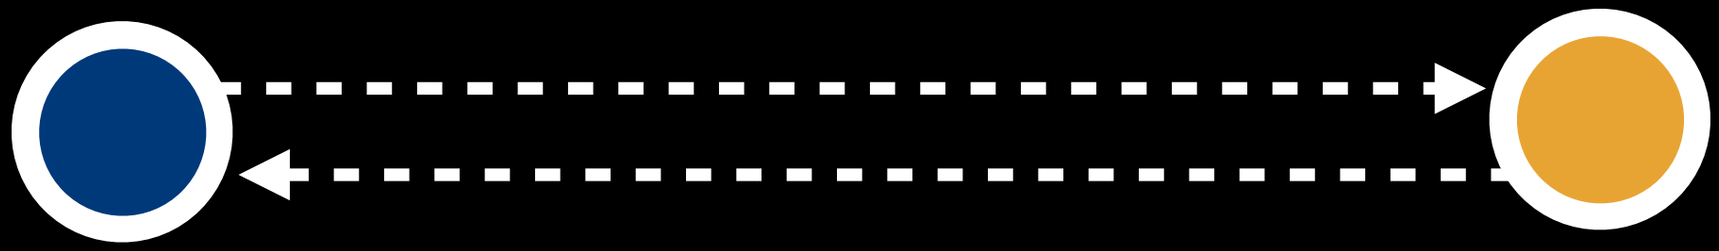

Pingers

- Send UDP and ICMP probes to target list
- Timestamp & Log results
- High ping-rate (up to 1Mpps)
- Set DSCP marking

Responders

- Receive/Reply to probe
- Timestamp
- Low load: thousands of pps
- Reflect DSCP value back

Why UDP?

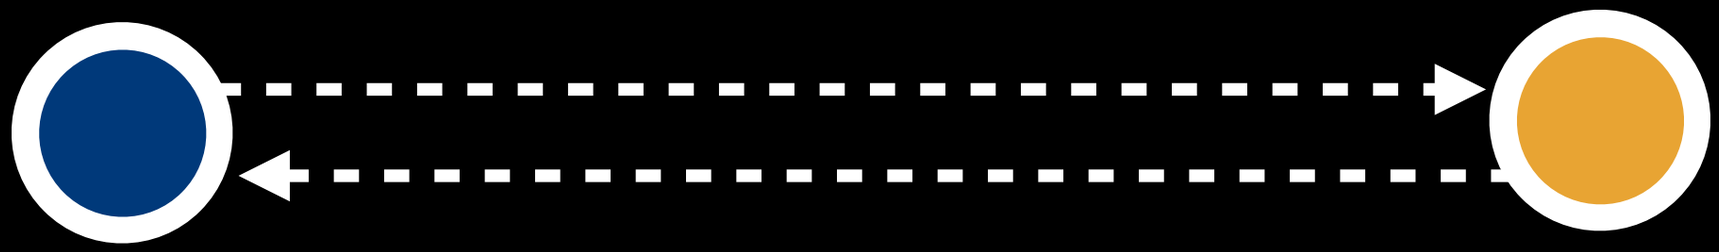

- No TCP RST packets
- Efficient ECMP coverage
- Extensible

Probe Structure

- Signature
- Send Time
- Receive Time
- Response Time
- Traffic Class

NetNORAD

Ping and Process Data

github.com/facebook/UdpPinger

Challenges

Tens of thousands of targets

Hundreds of pingers

Lots of data to process

We really do not care about each host…
…The unit of interest is cluster health

Pinging inside clusters

- Detect issues with rack switches
- Dedicated pingers per cluster
- Probe ALL machines in cluster
- Store time-series per host/rack
- Lags real-time by 2 minutes

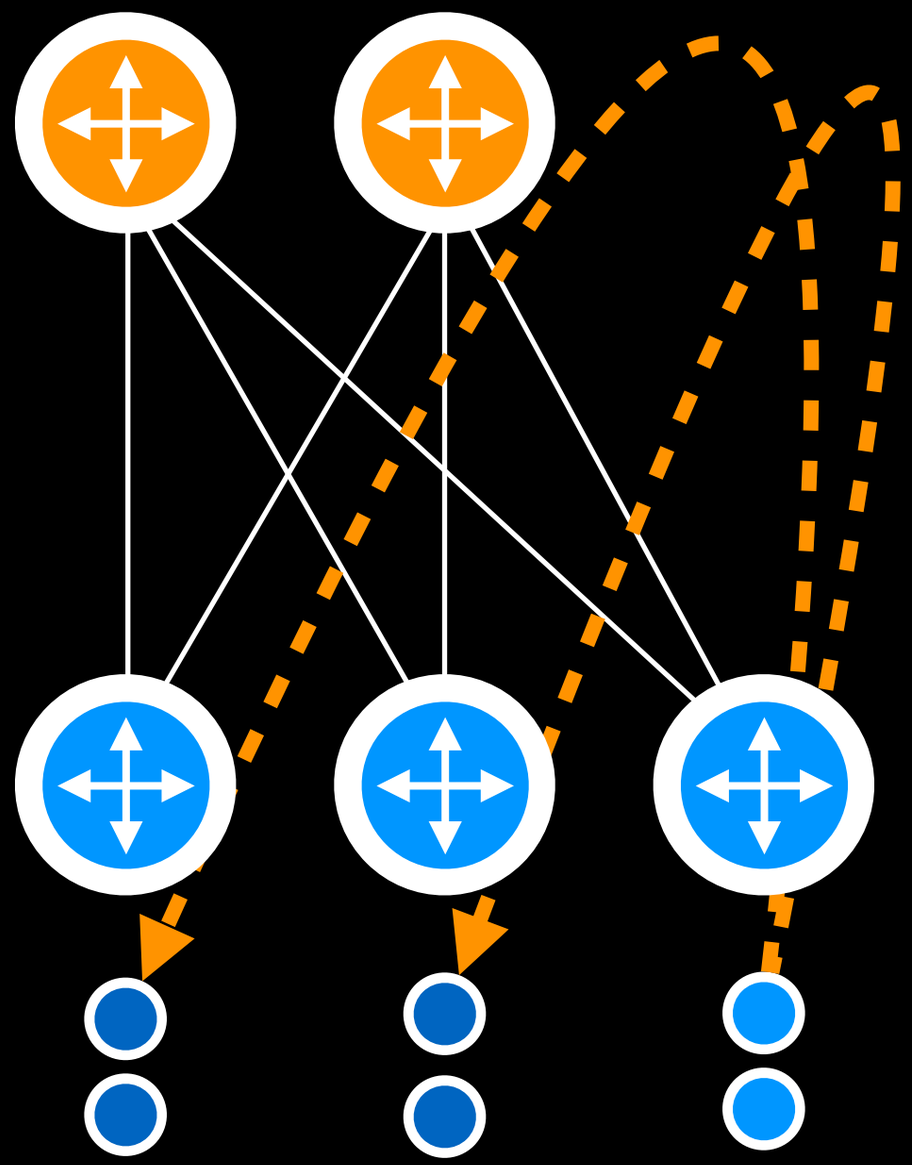

Pinging the clusters

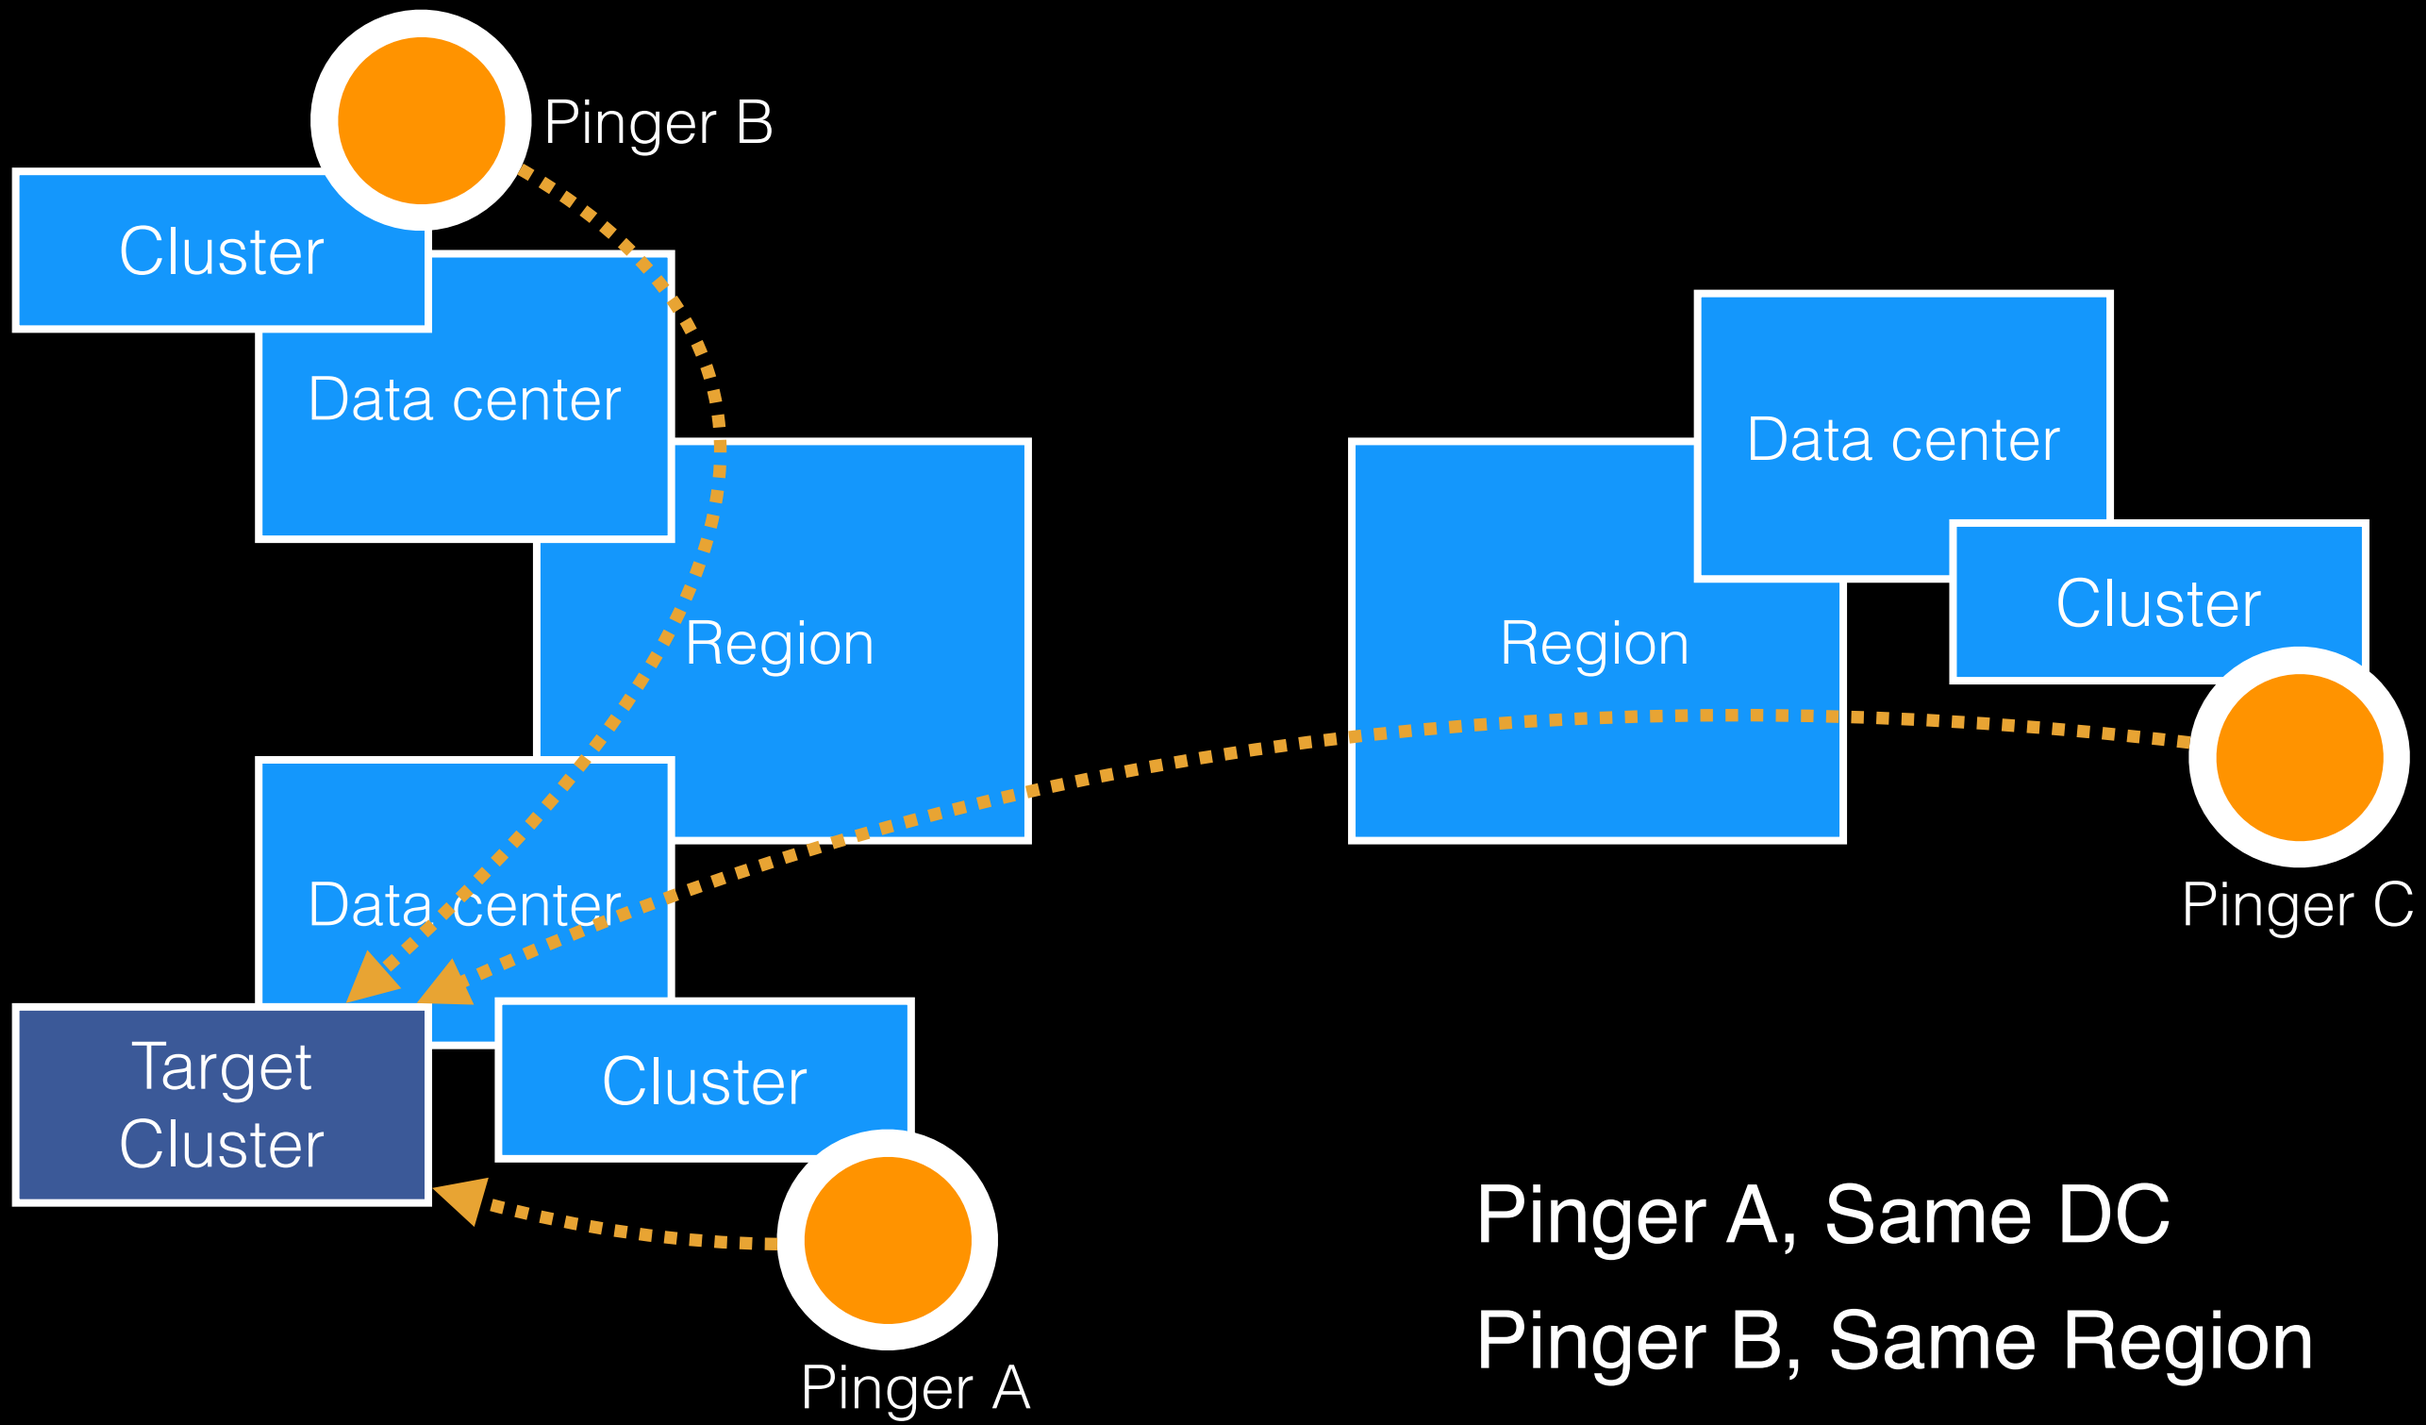

Pinger A, Same DC

Pinger B, Same Region

Pinger C, Outside of region

Alarming on Loss

- Build packet loss time-series
- Track percentiles
- Time to detect loss? 20 Seconds

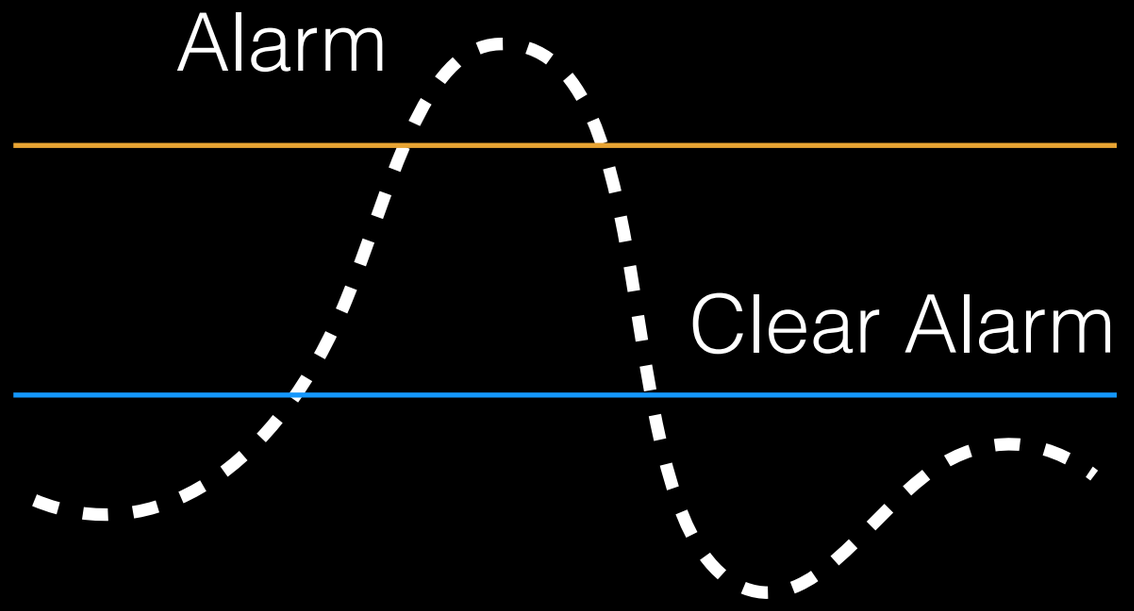

Finding the problem

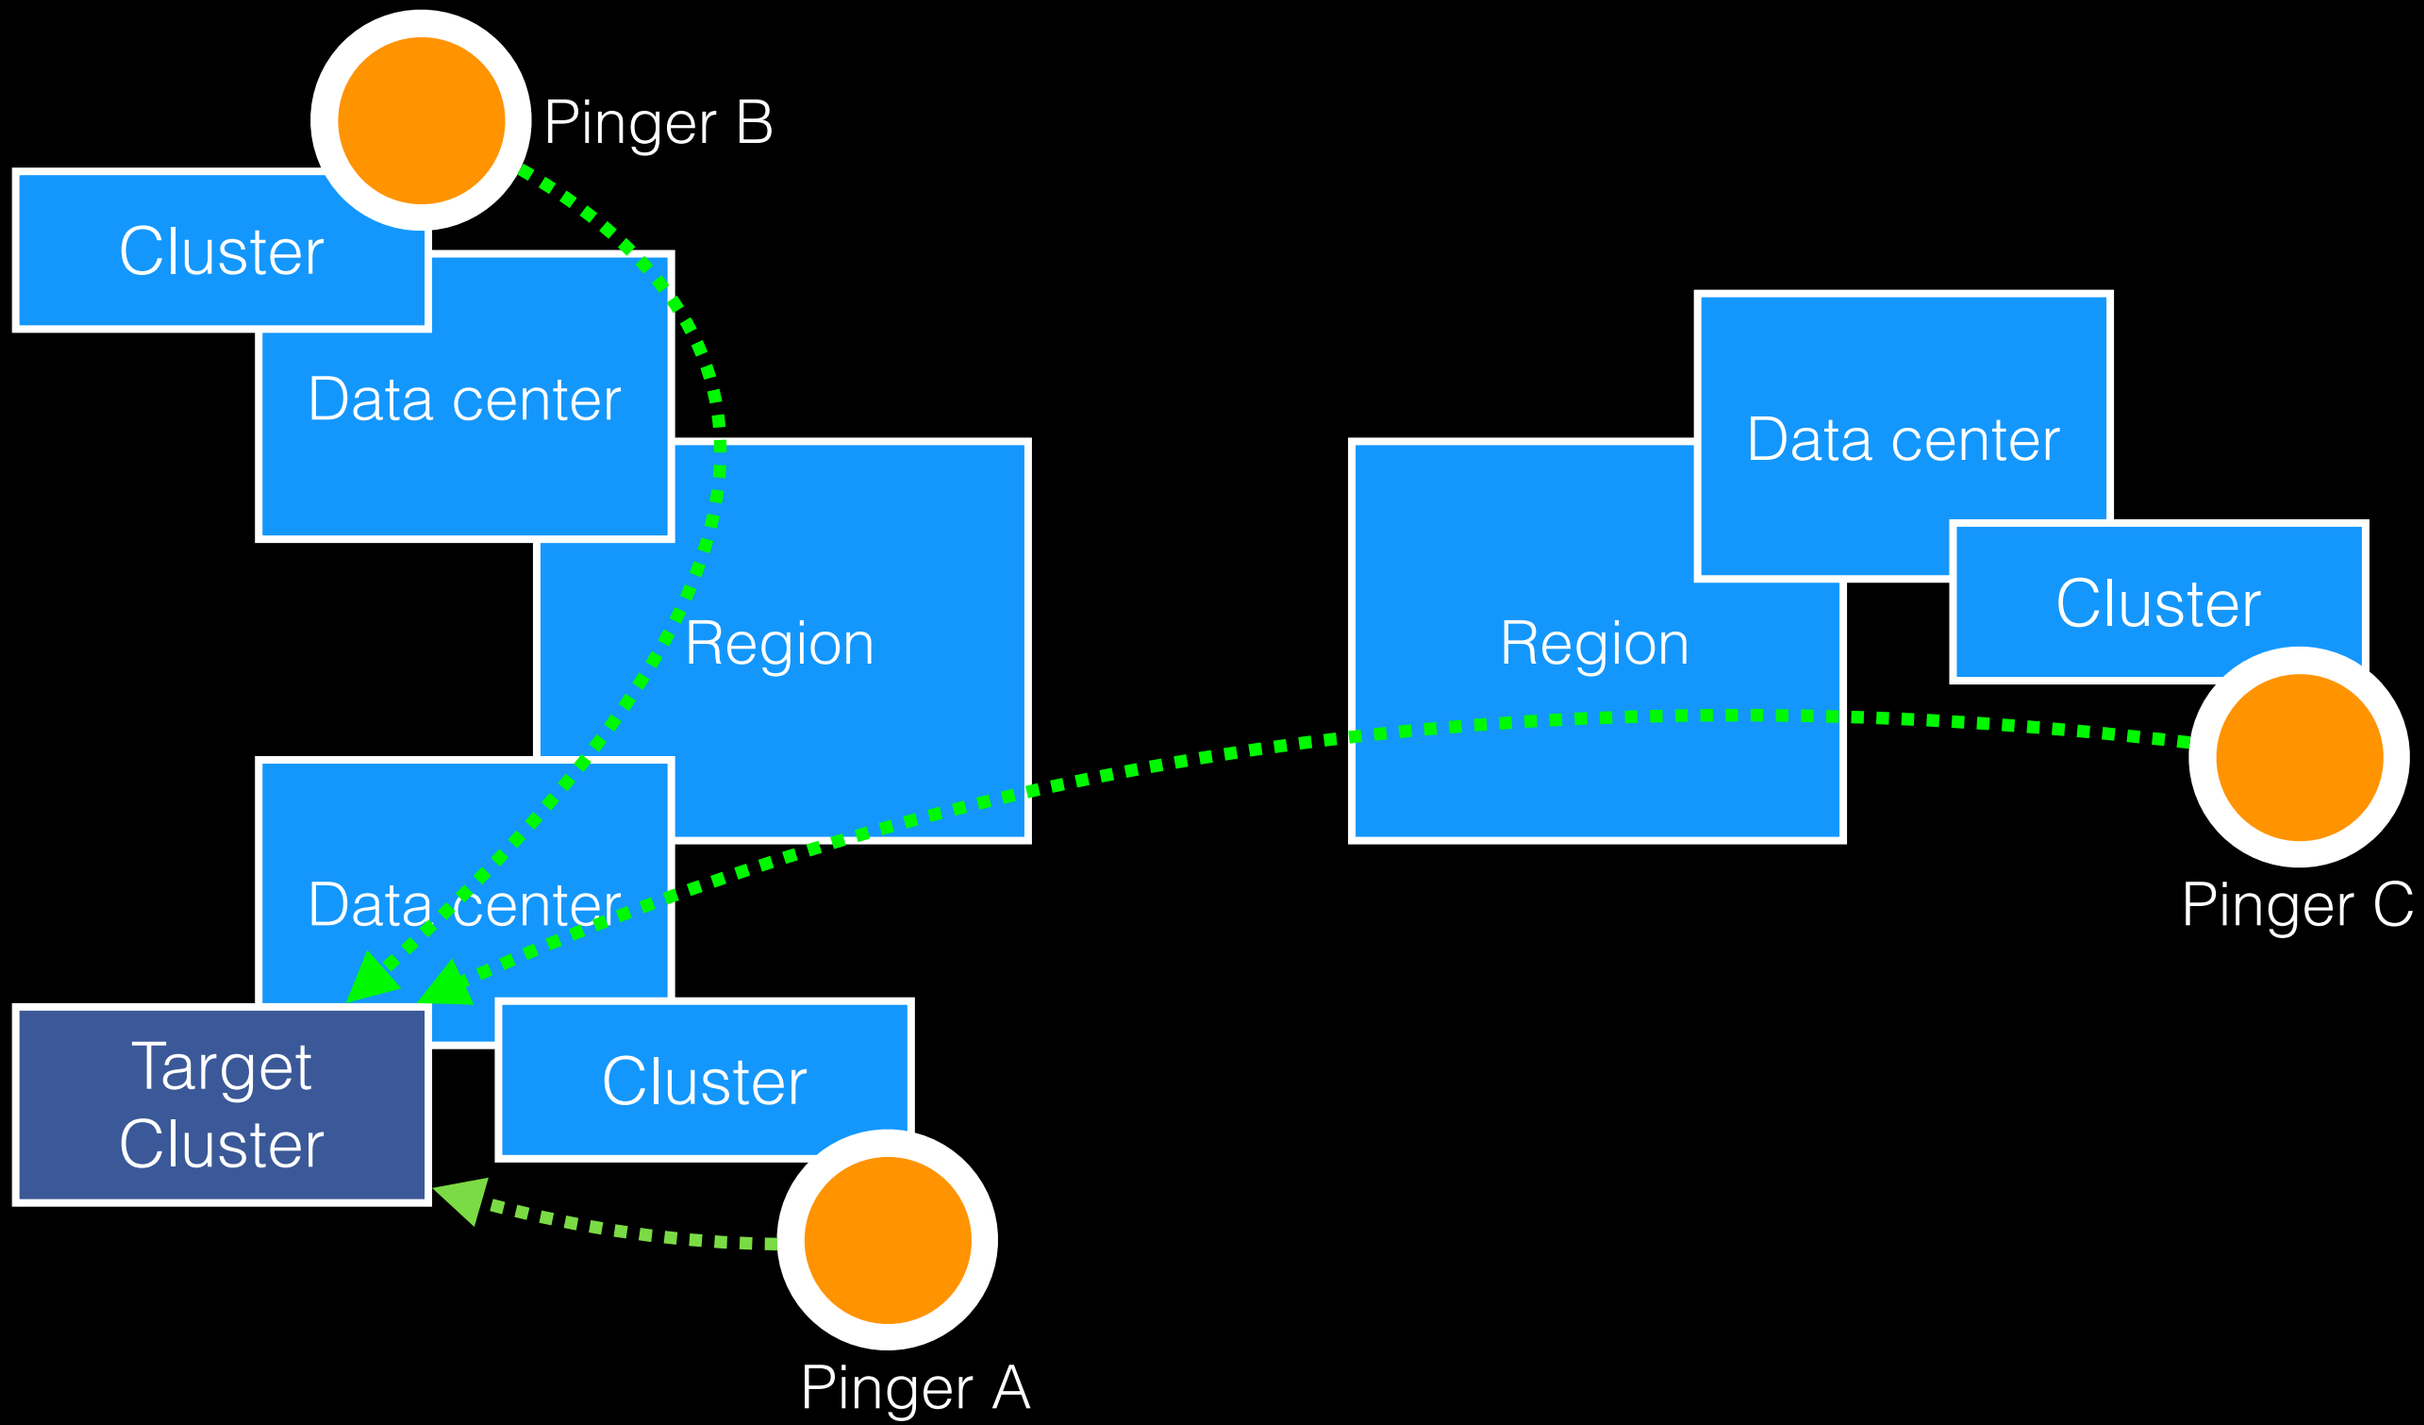

Finding the problem

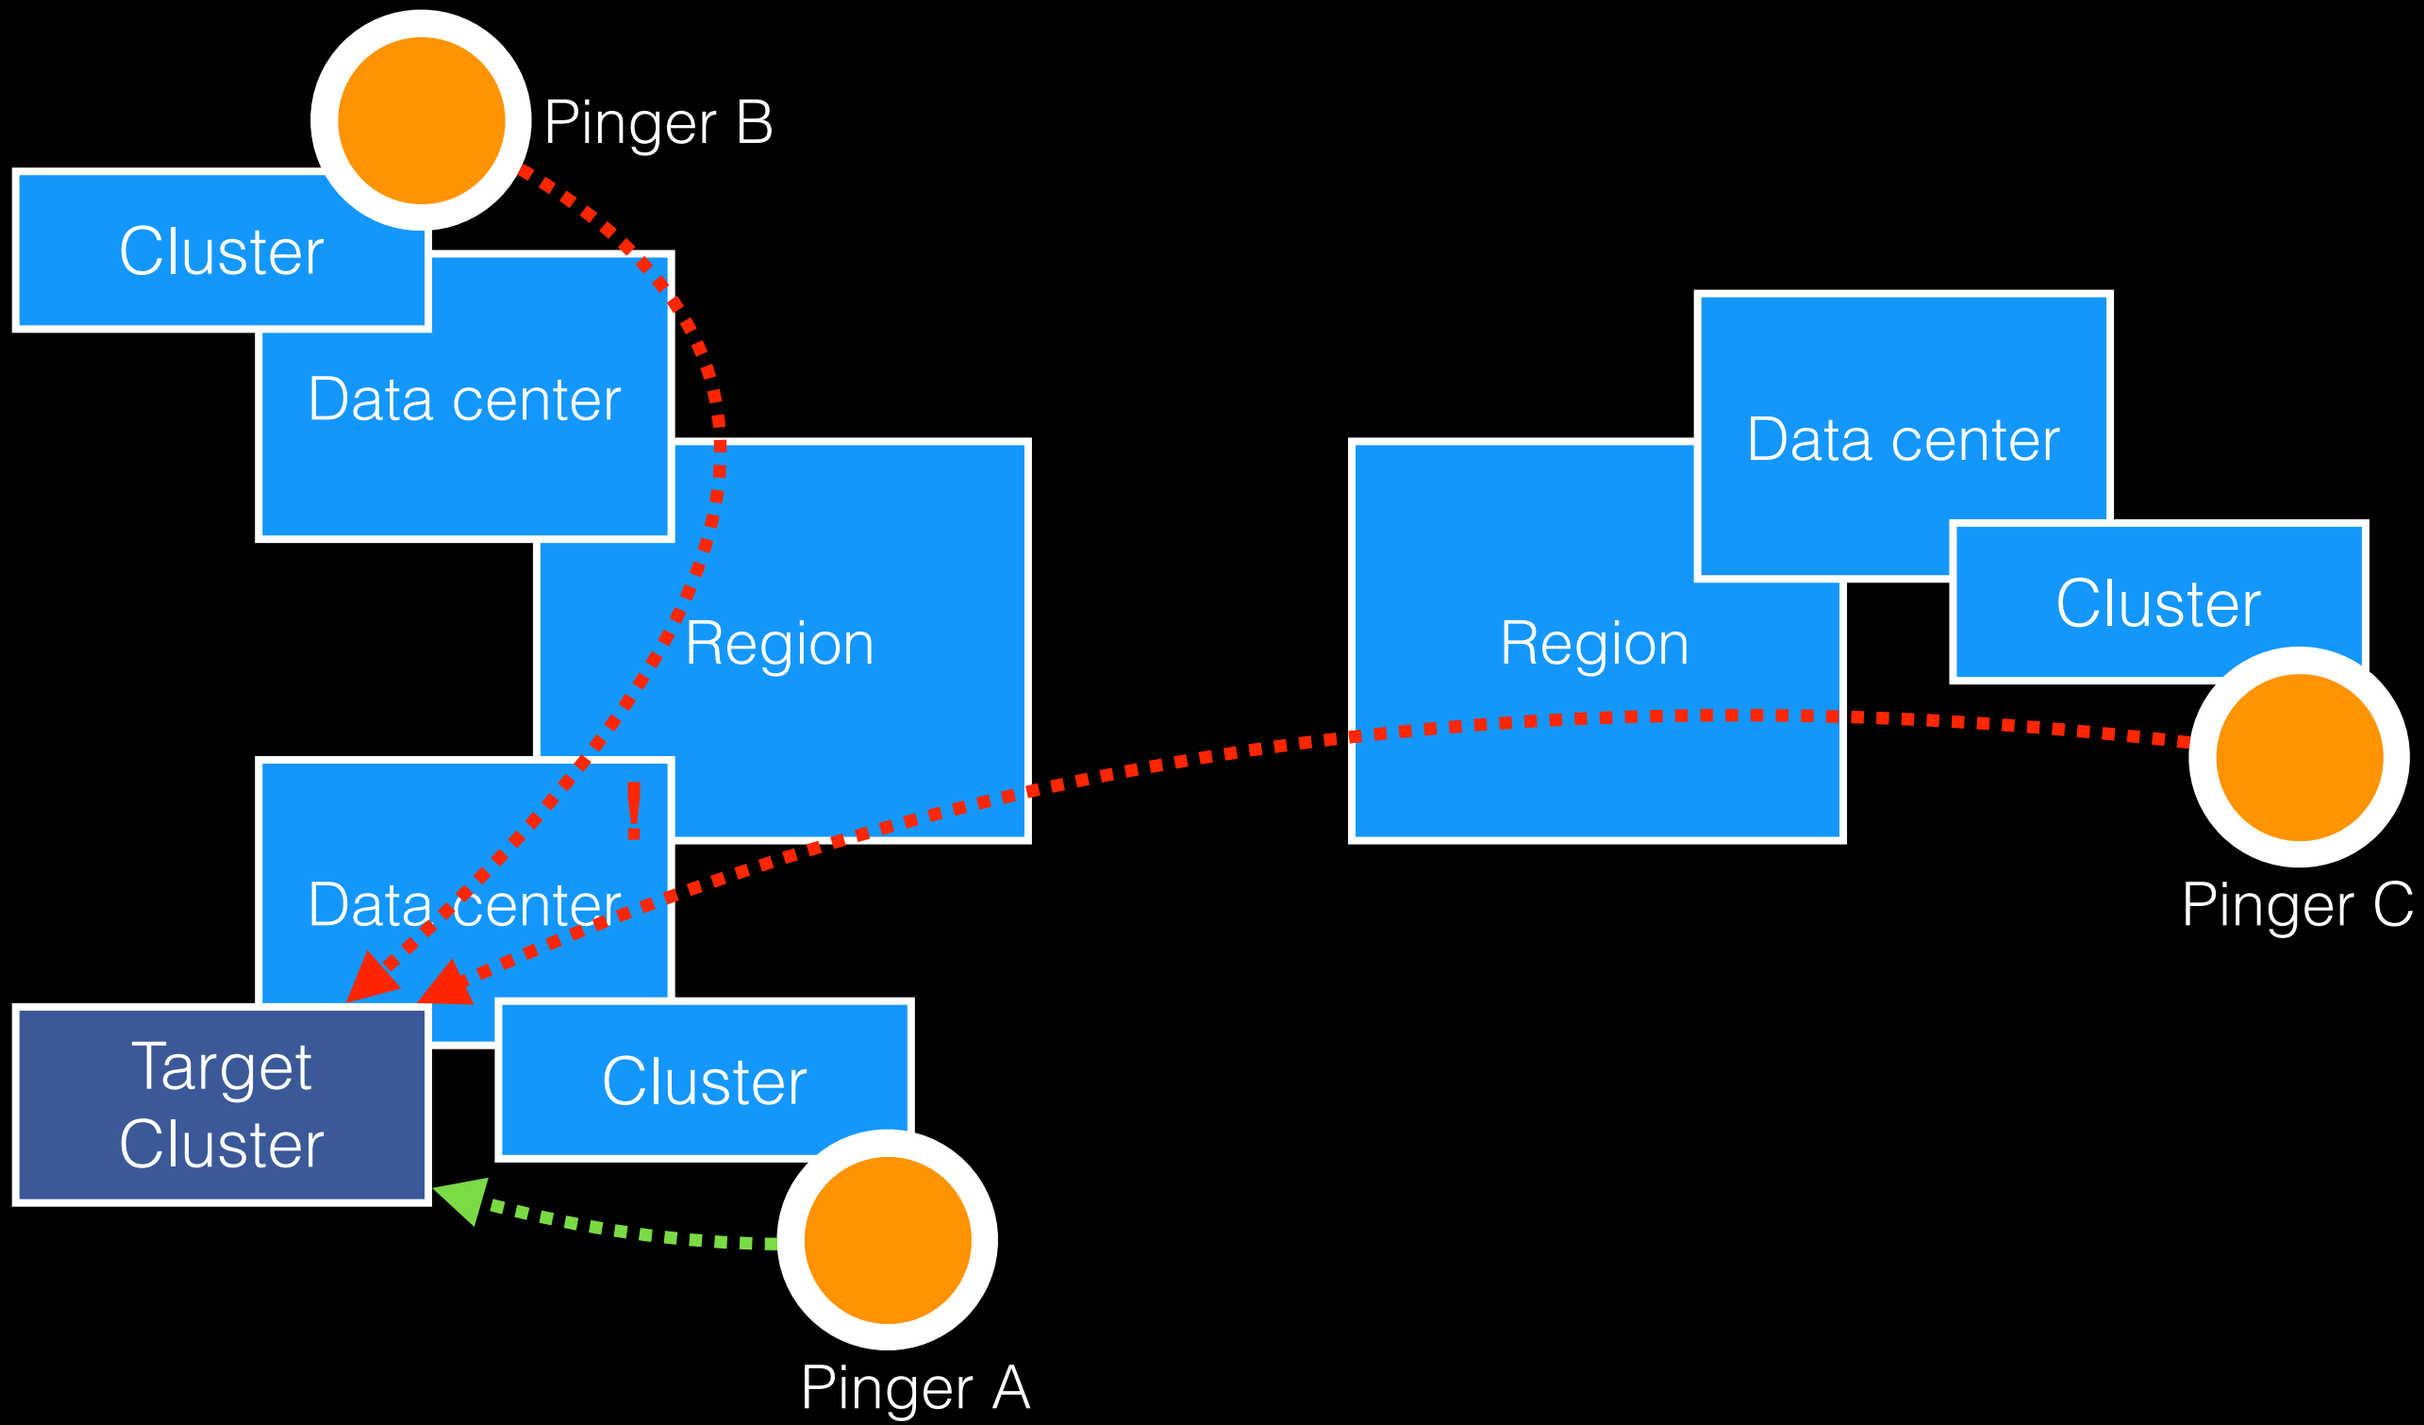

Finding the problem

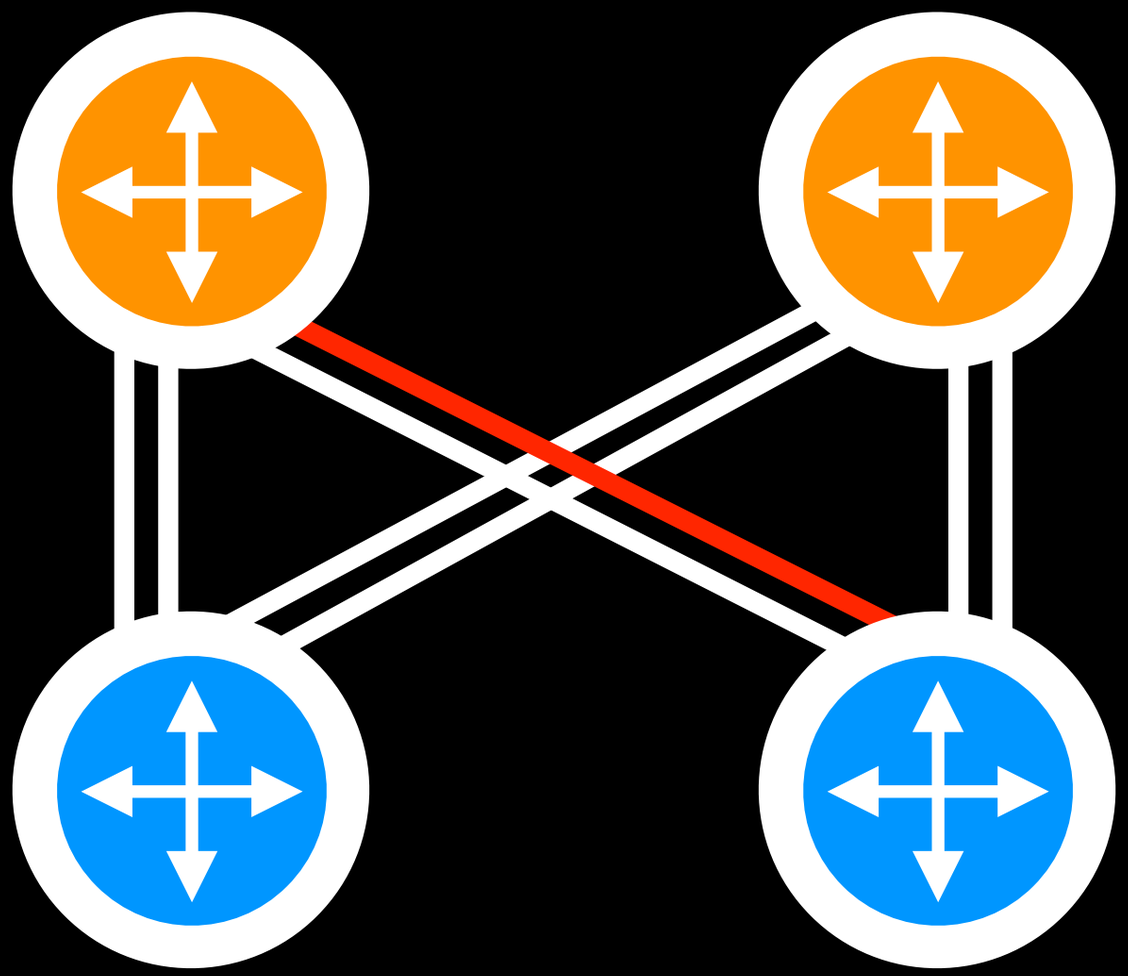

Each layer of the network contains many devices
...and many more links

fbtracert

Isolating Network Faults

github.com/facebook/fbtracert

fbtracert

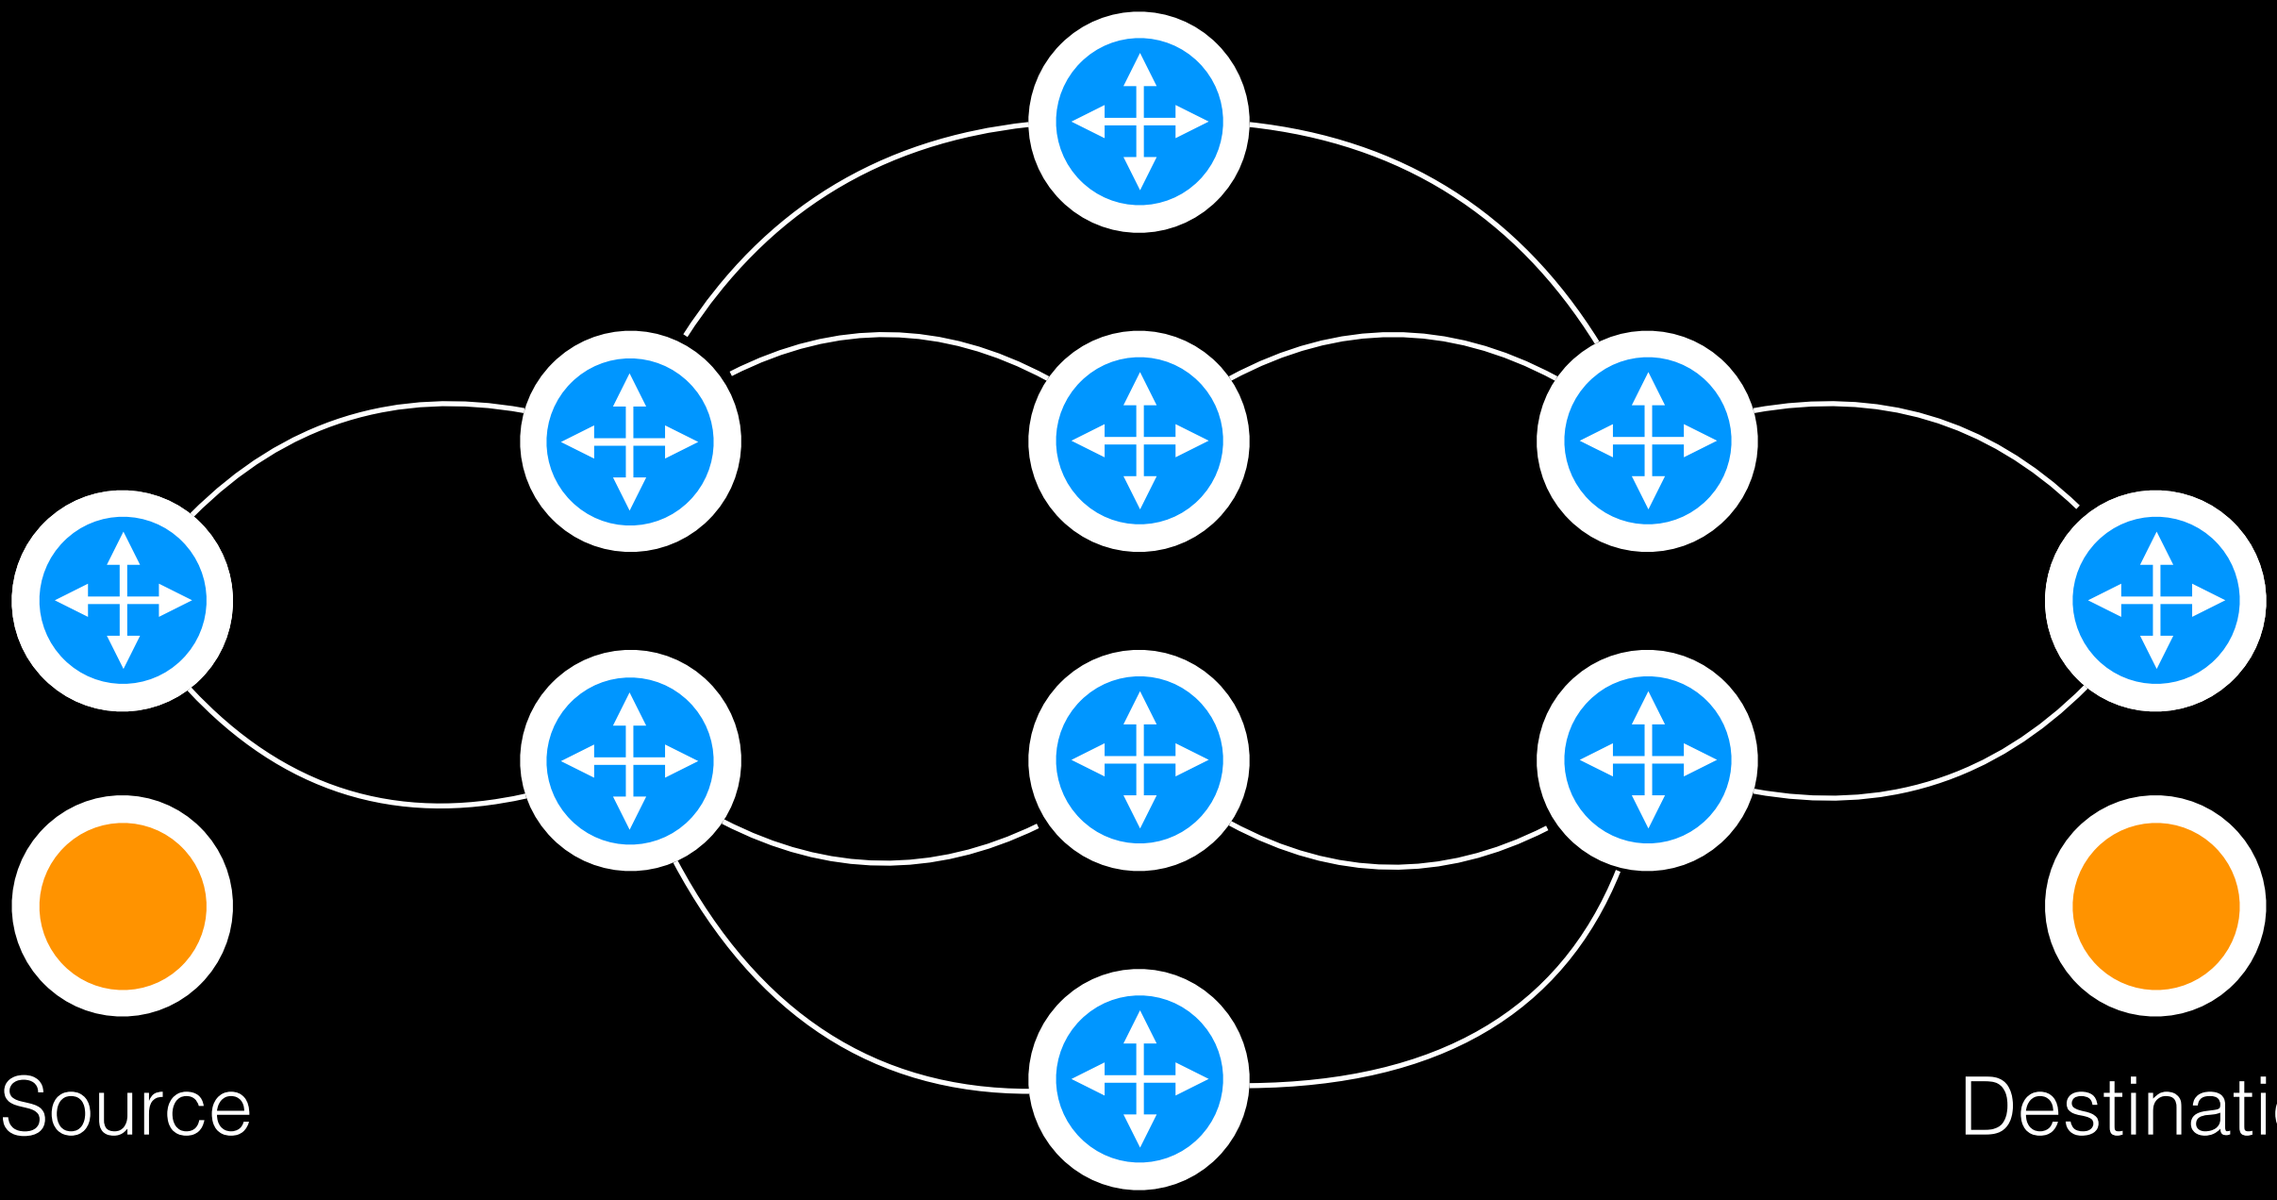

fbtracert

Src Port
13370
13371
13372
13373

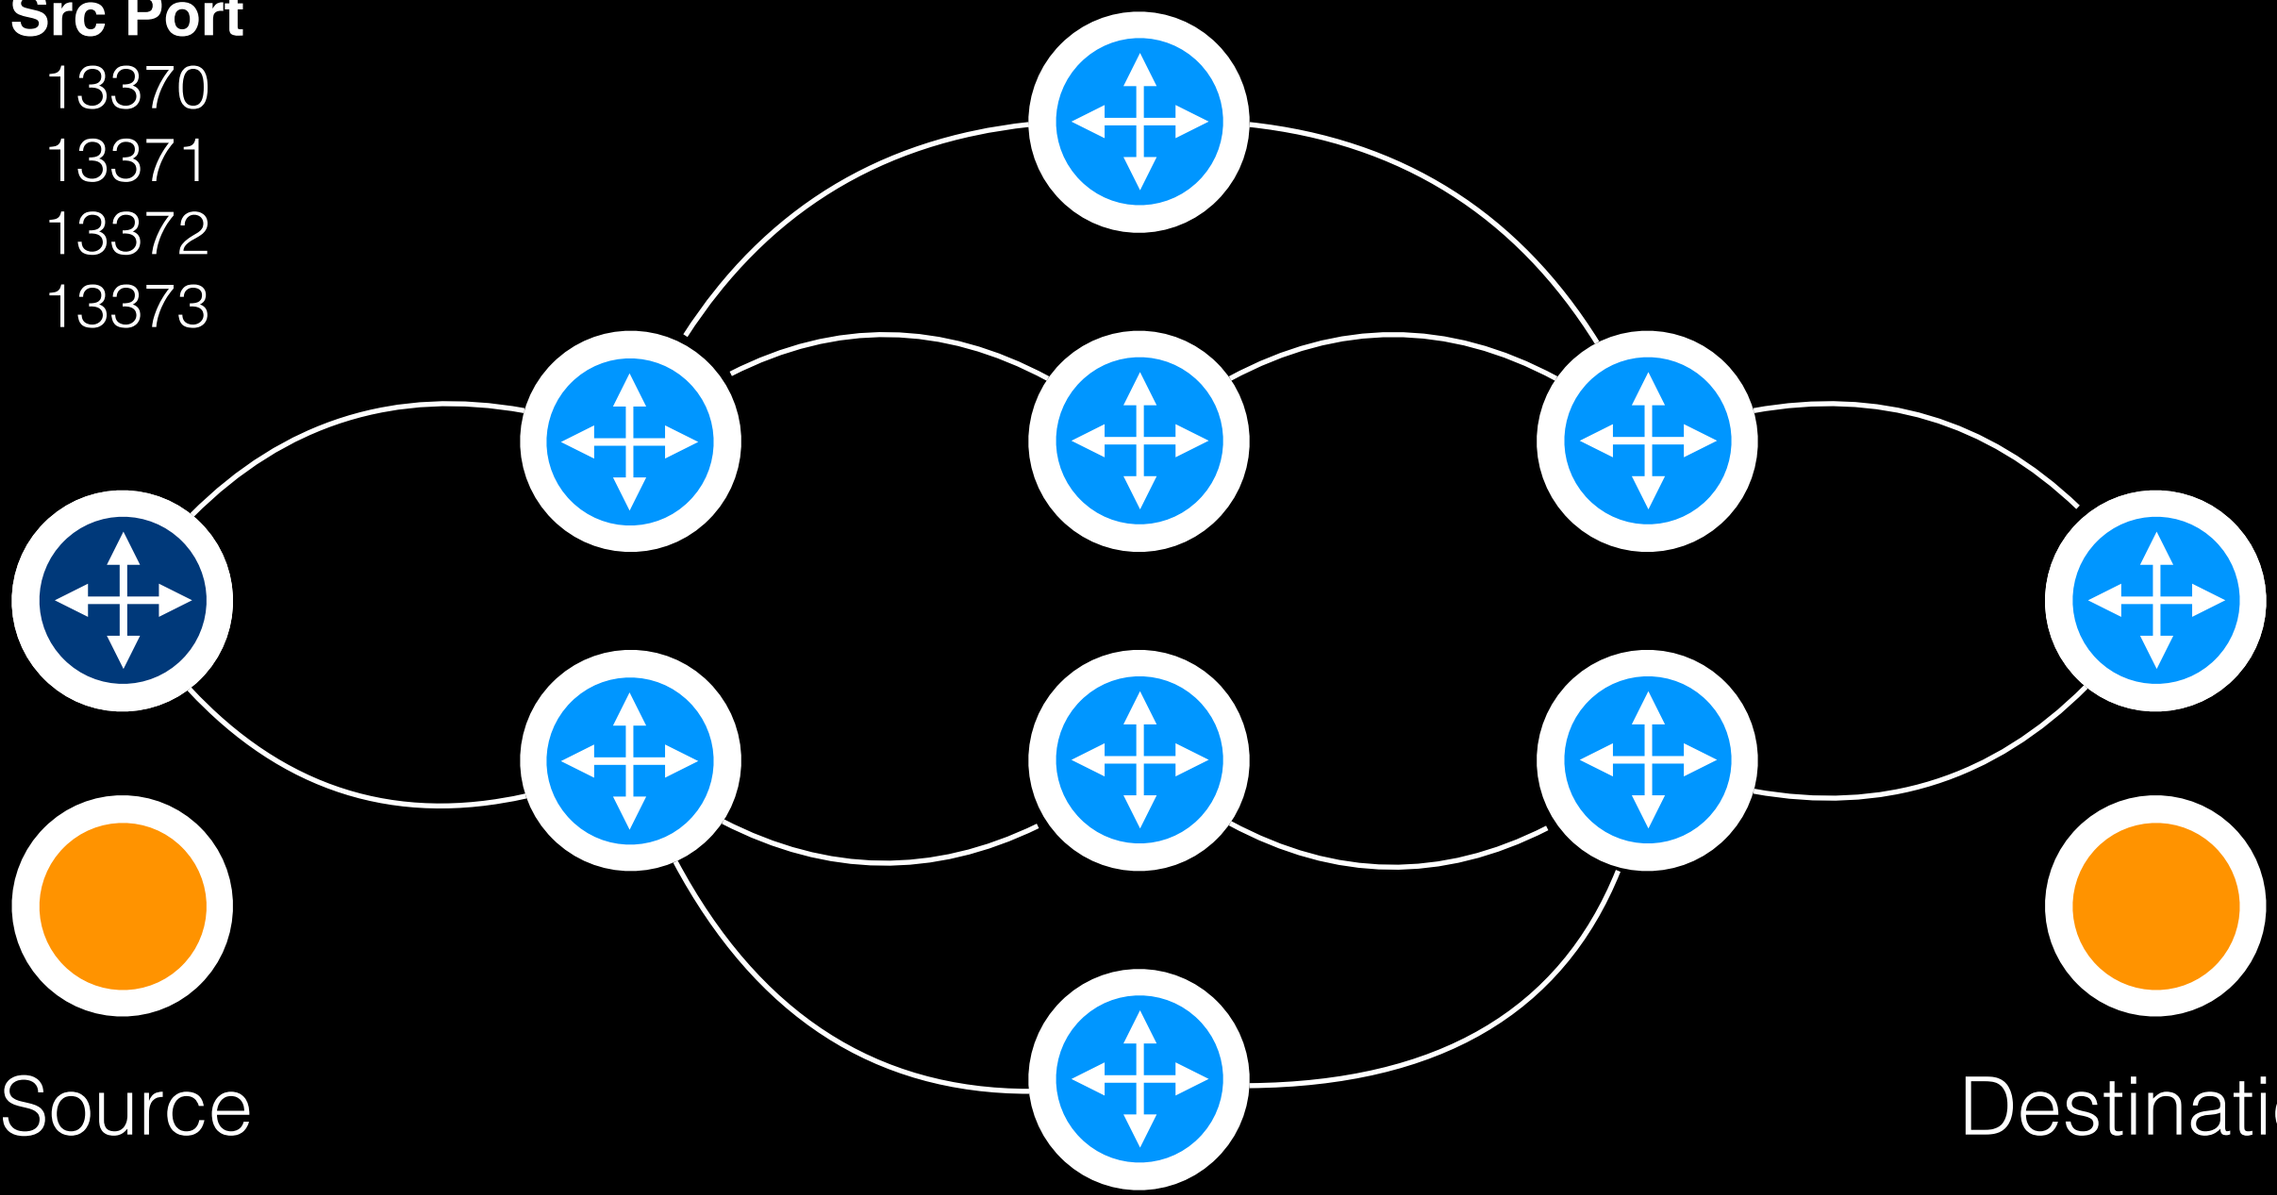

Source

Destination

fbtracert

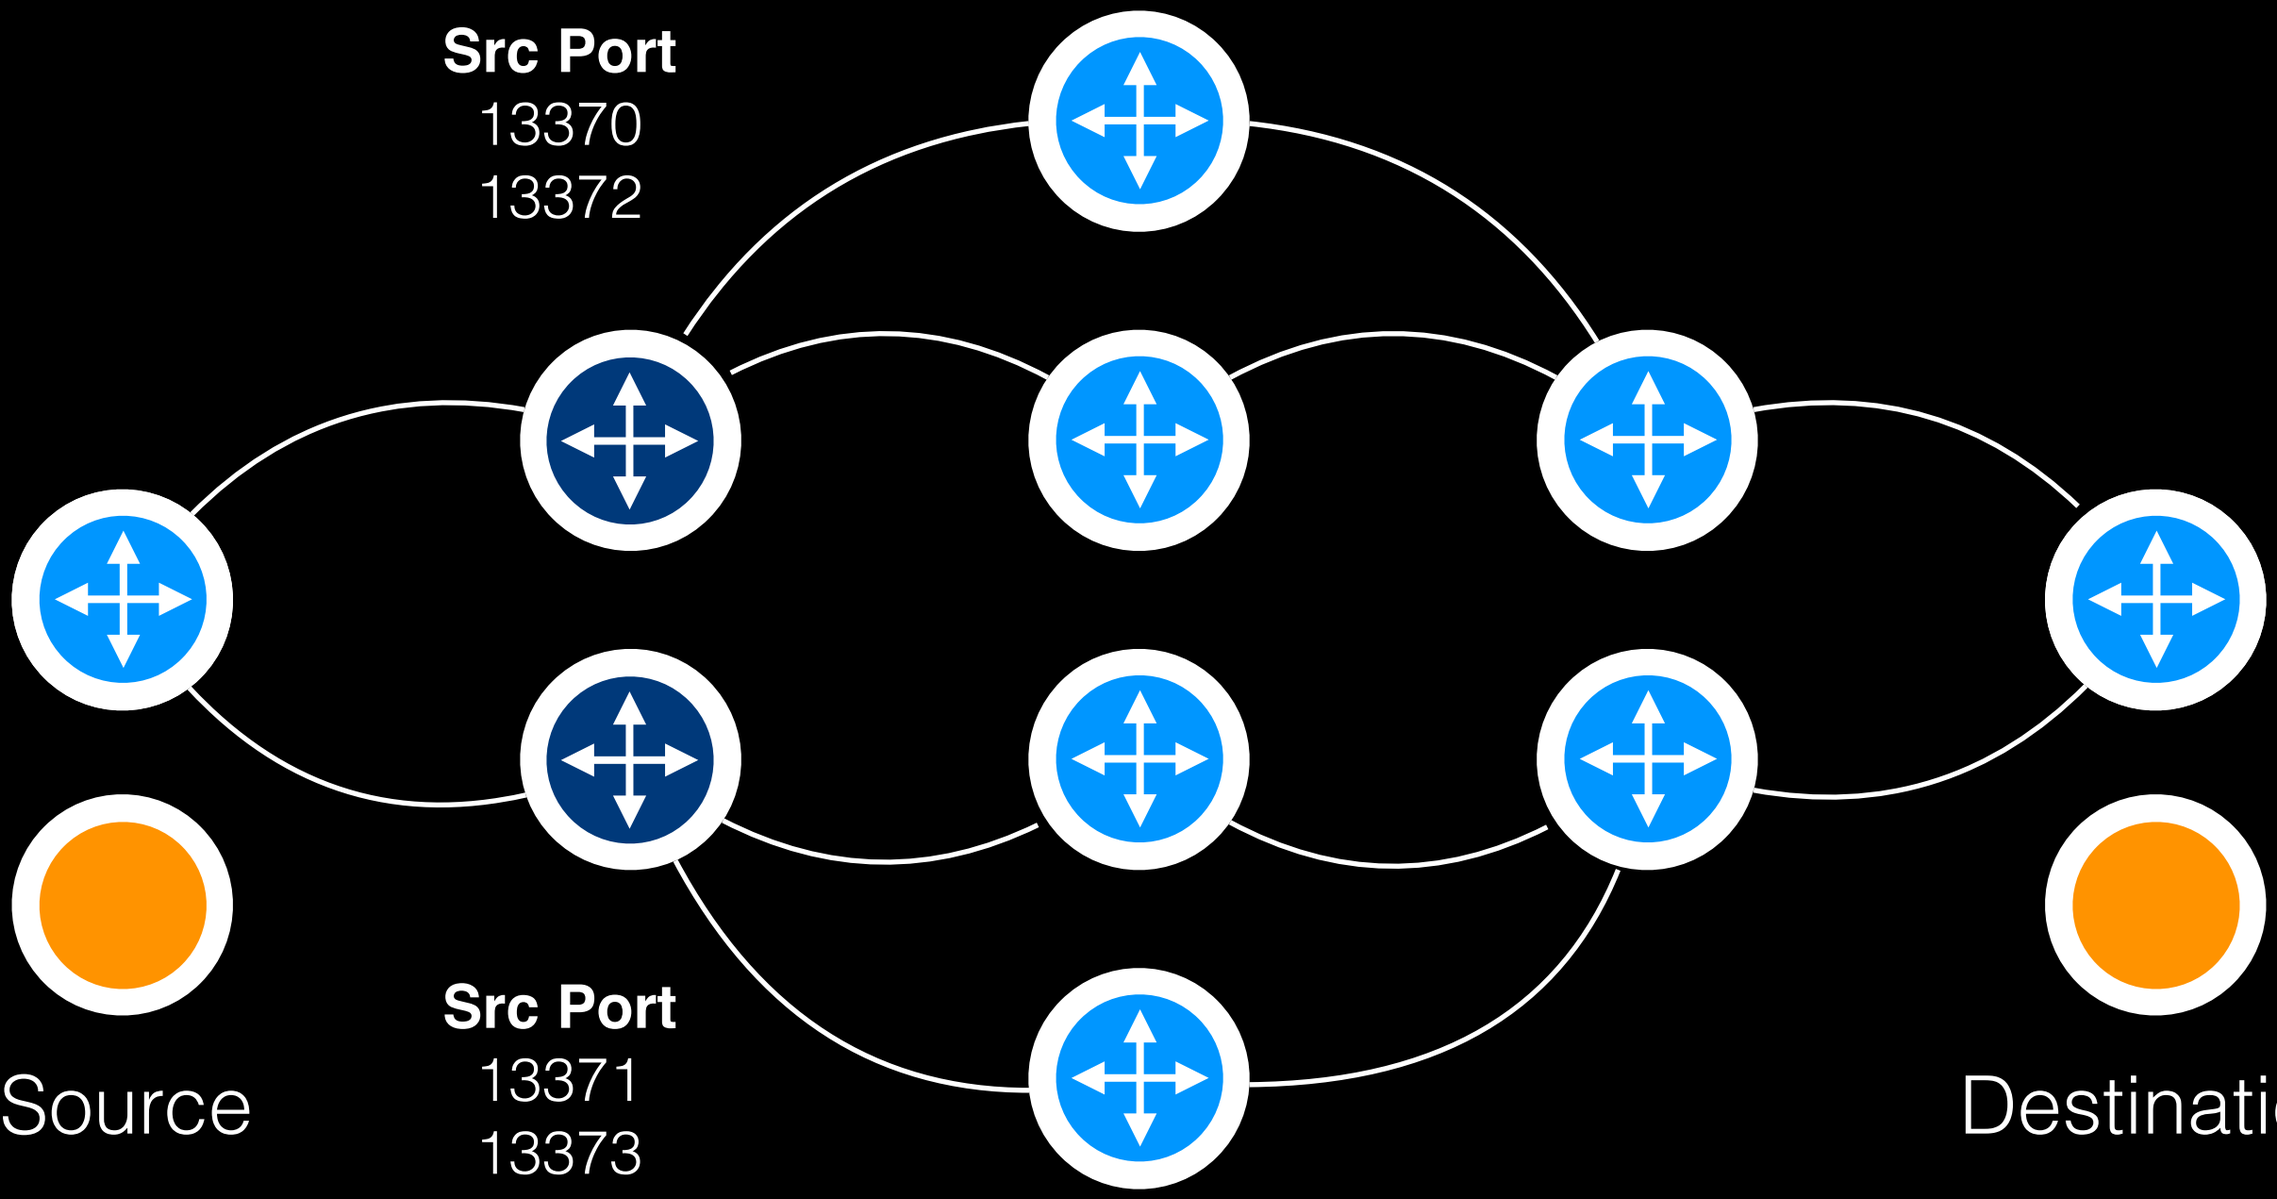
Src Port
13370
13372
Src Port
13371
13373
Source
Destination

fbtracert

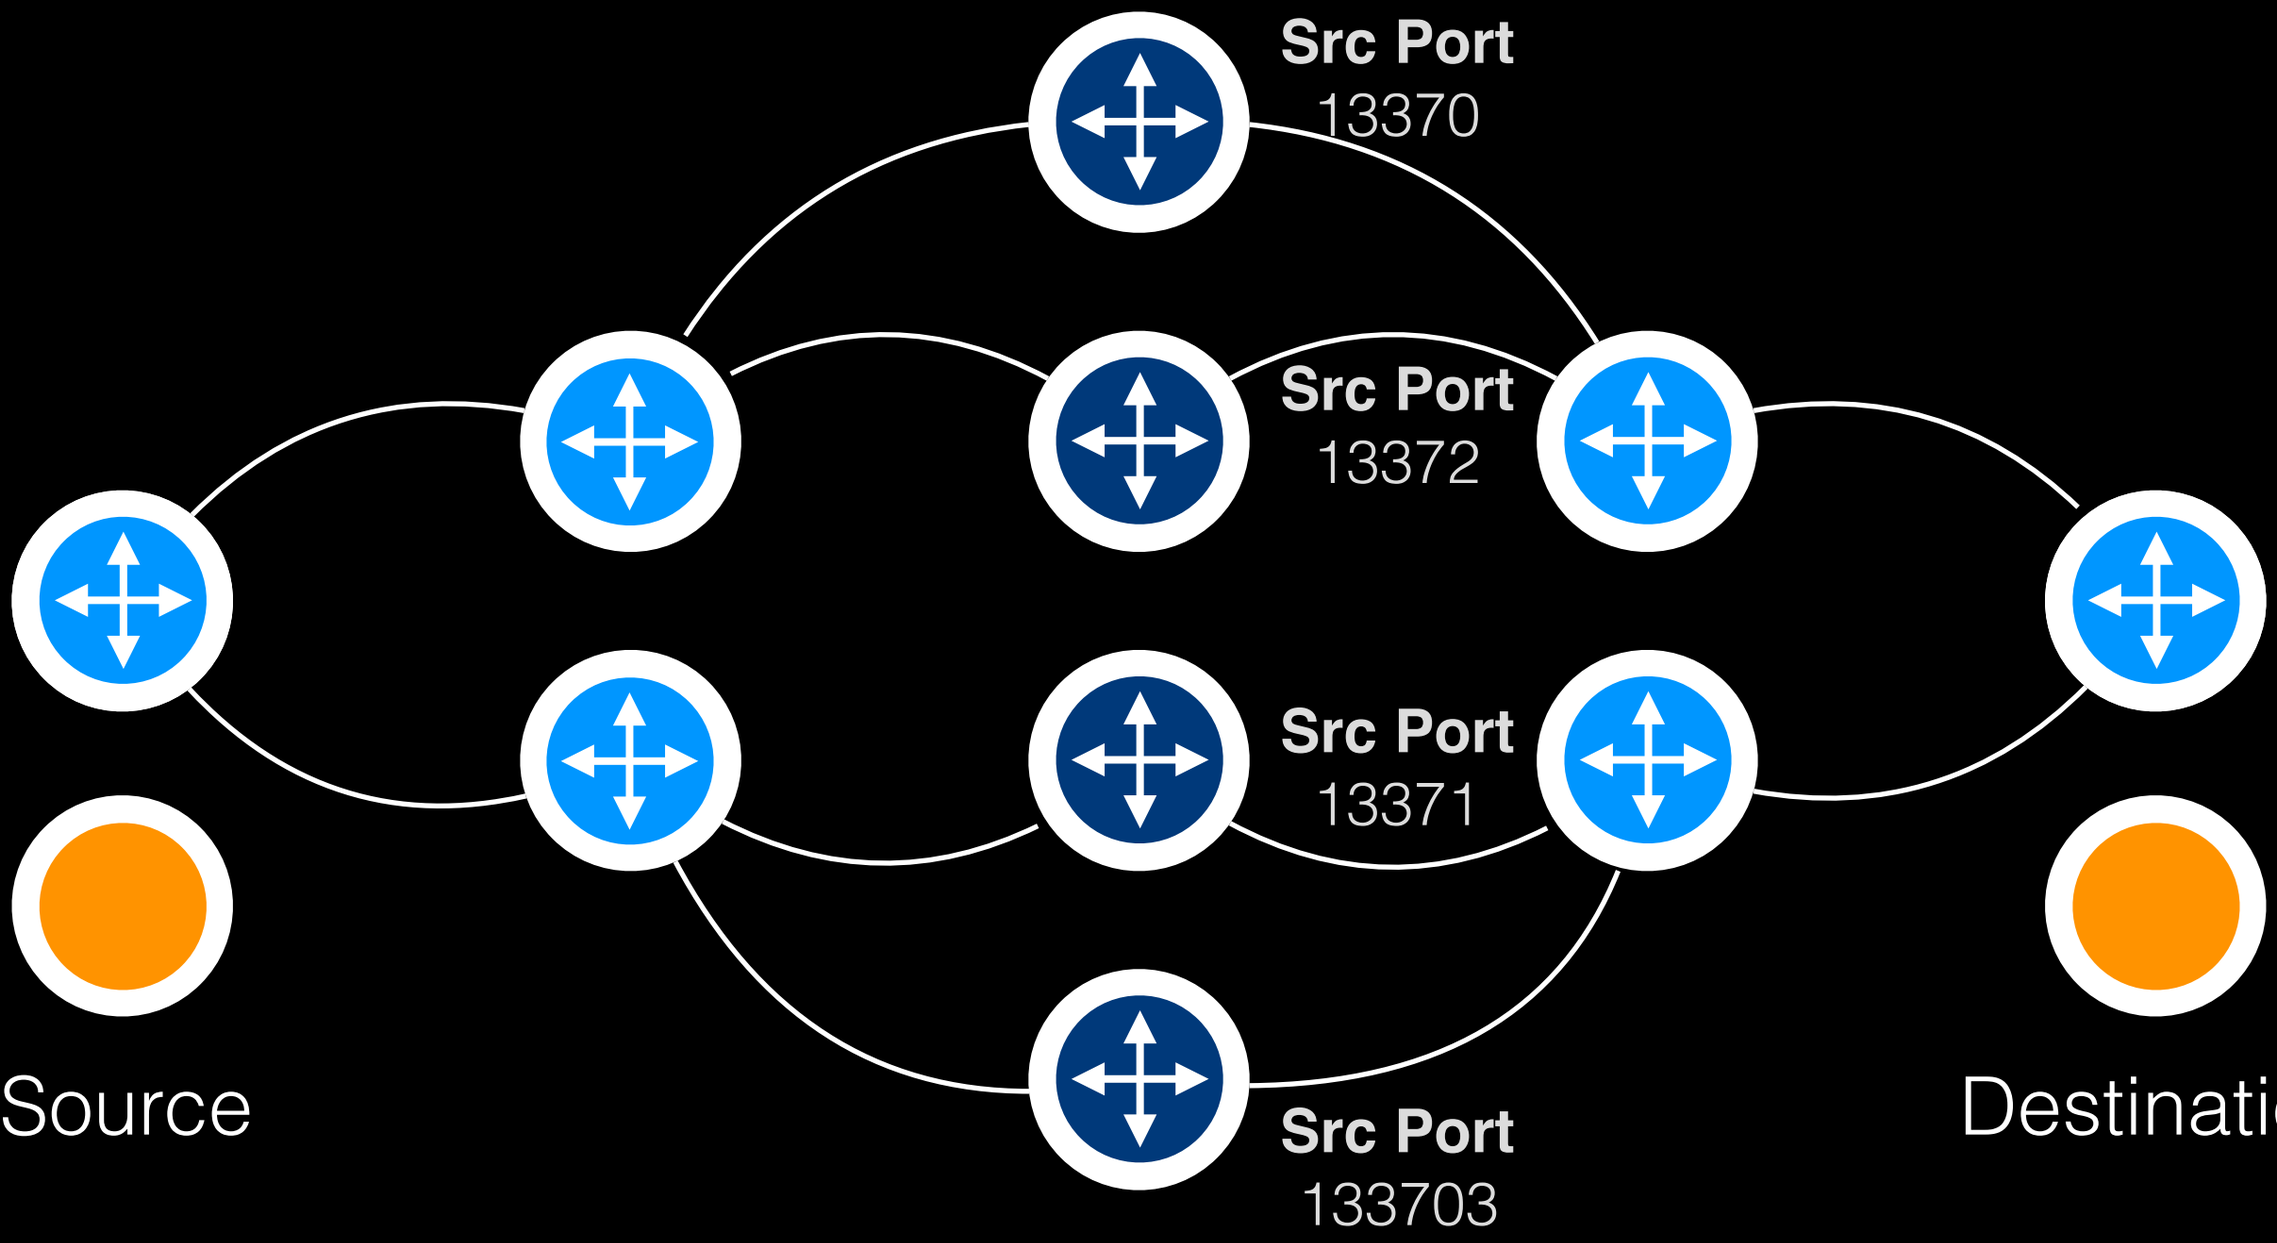
Src Port
13370
Src Port
13372
Src Port
13371
Src Port
133703
Source
Destination

fbtracert

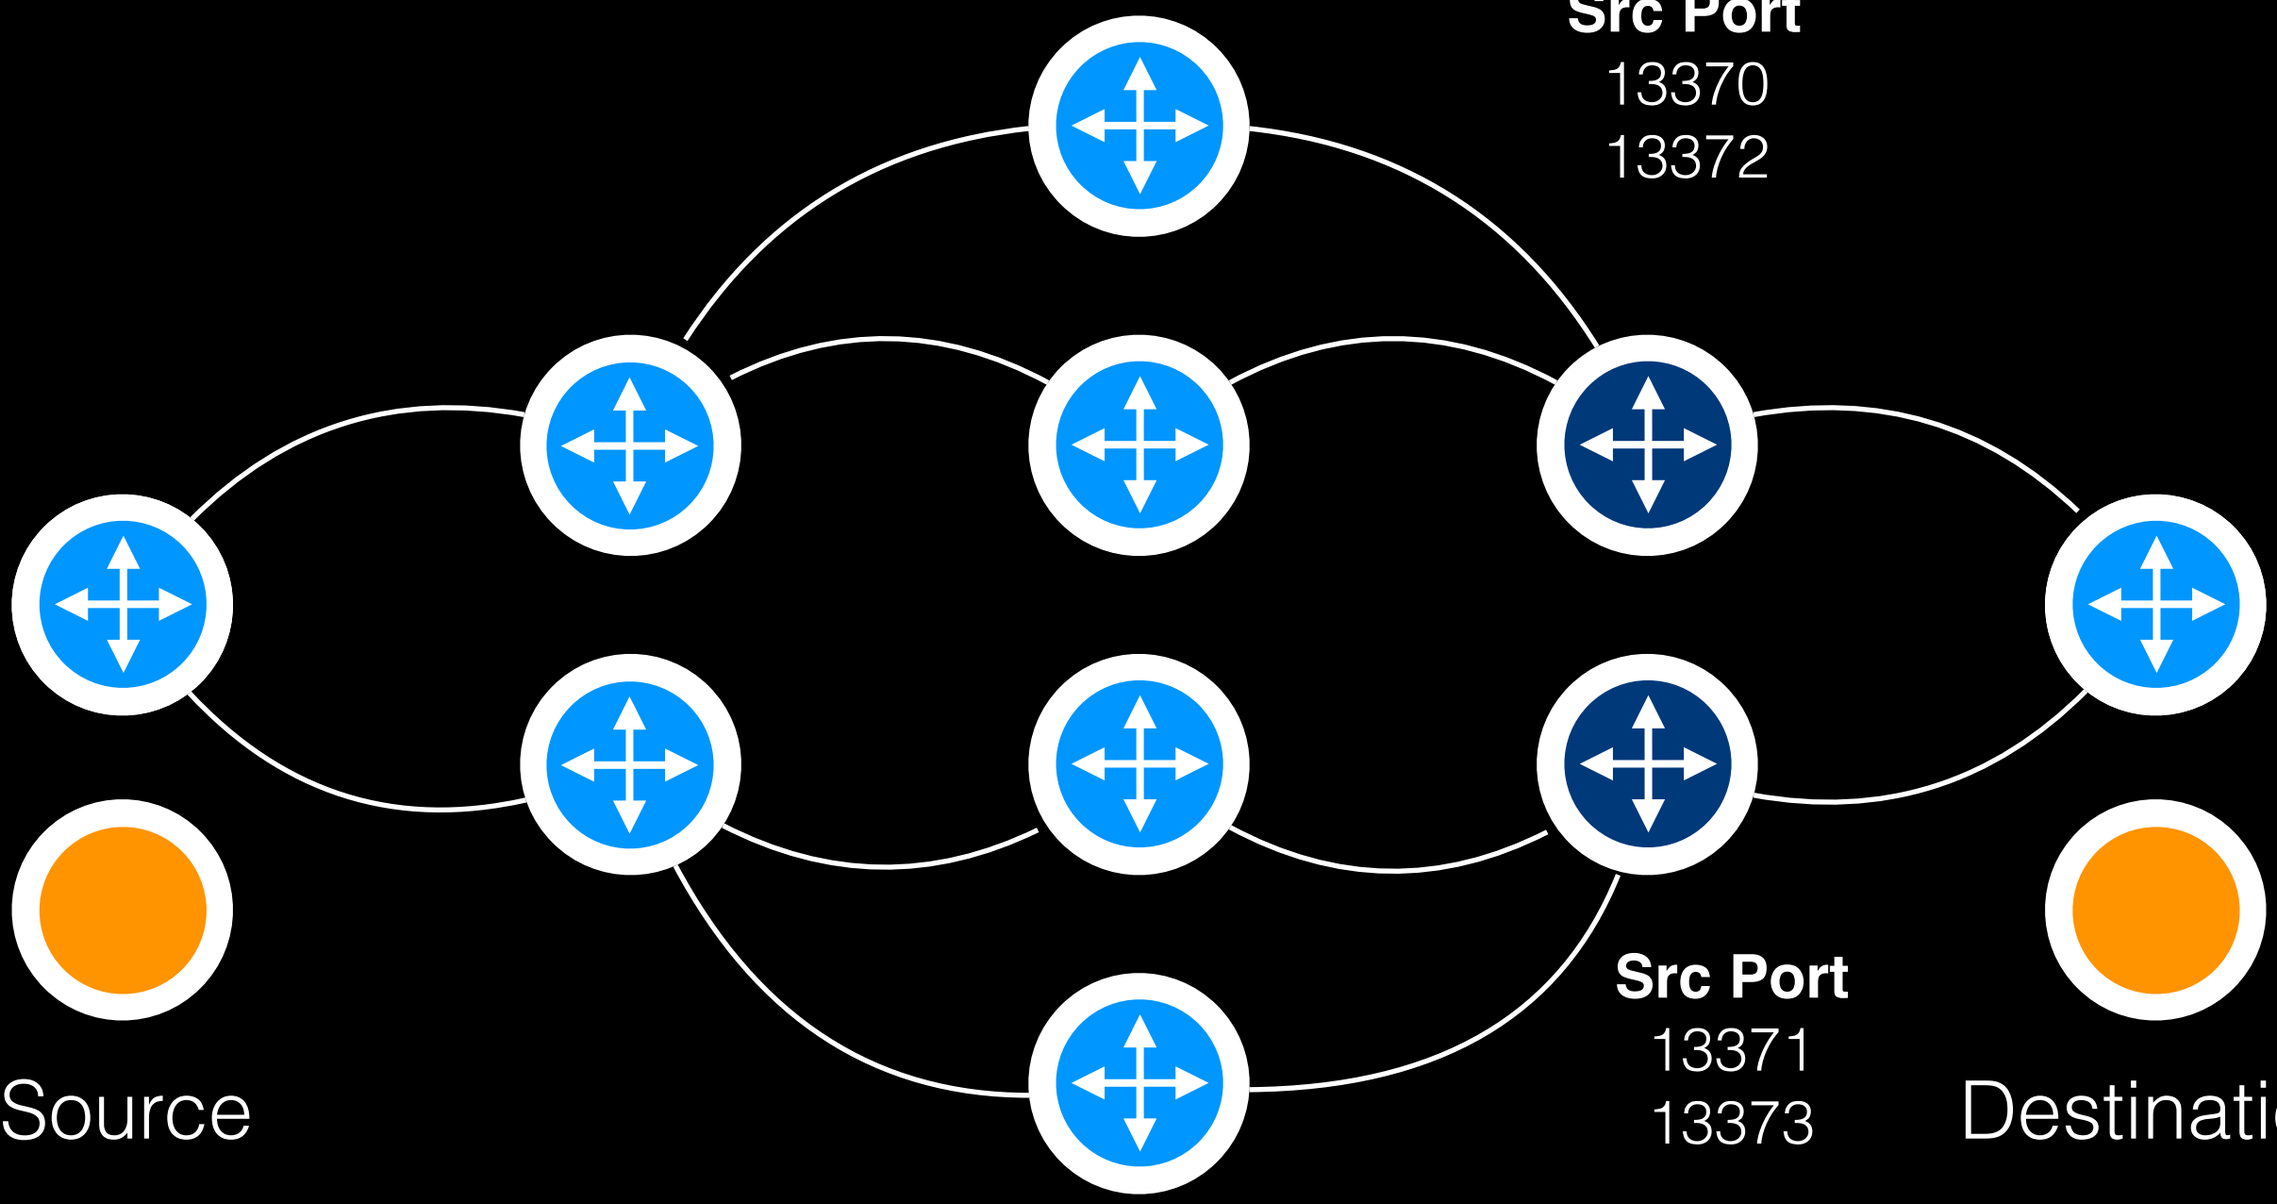
Src Port
13370
13372
Src Port
13371
13373
Source
Destination

fbtracert

Src Port
13370
13371
13372
13373

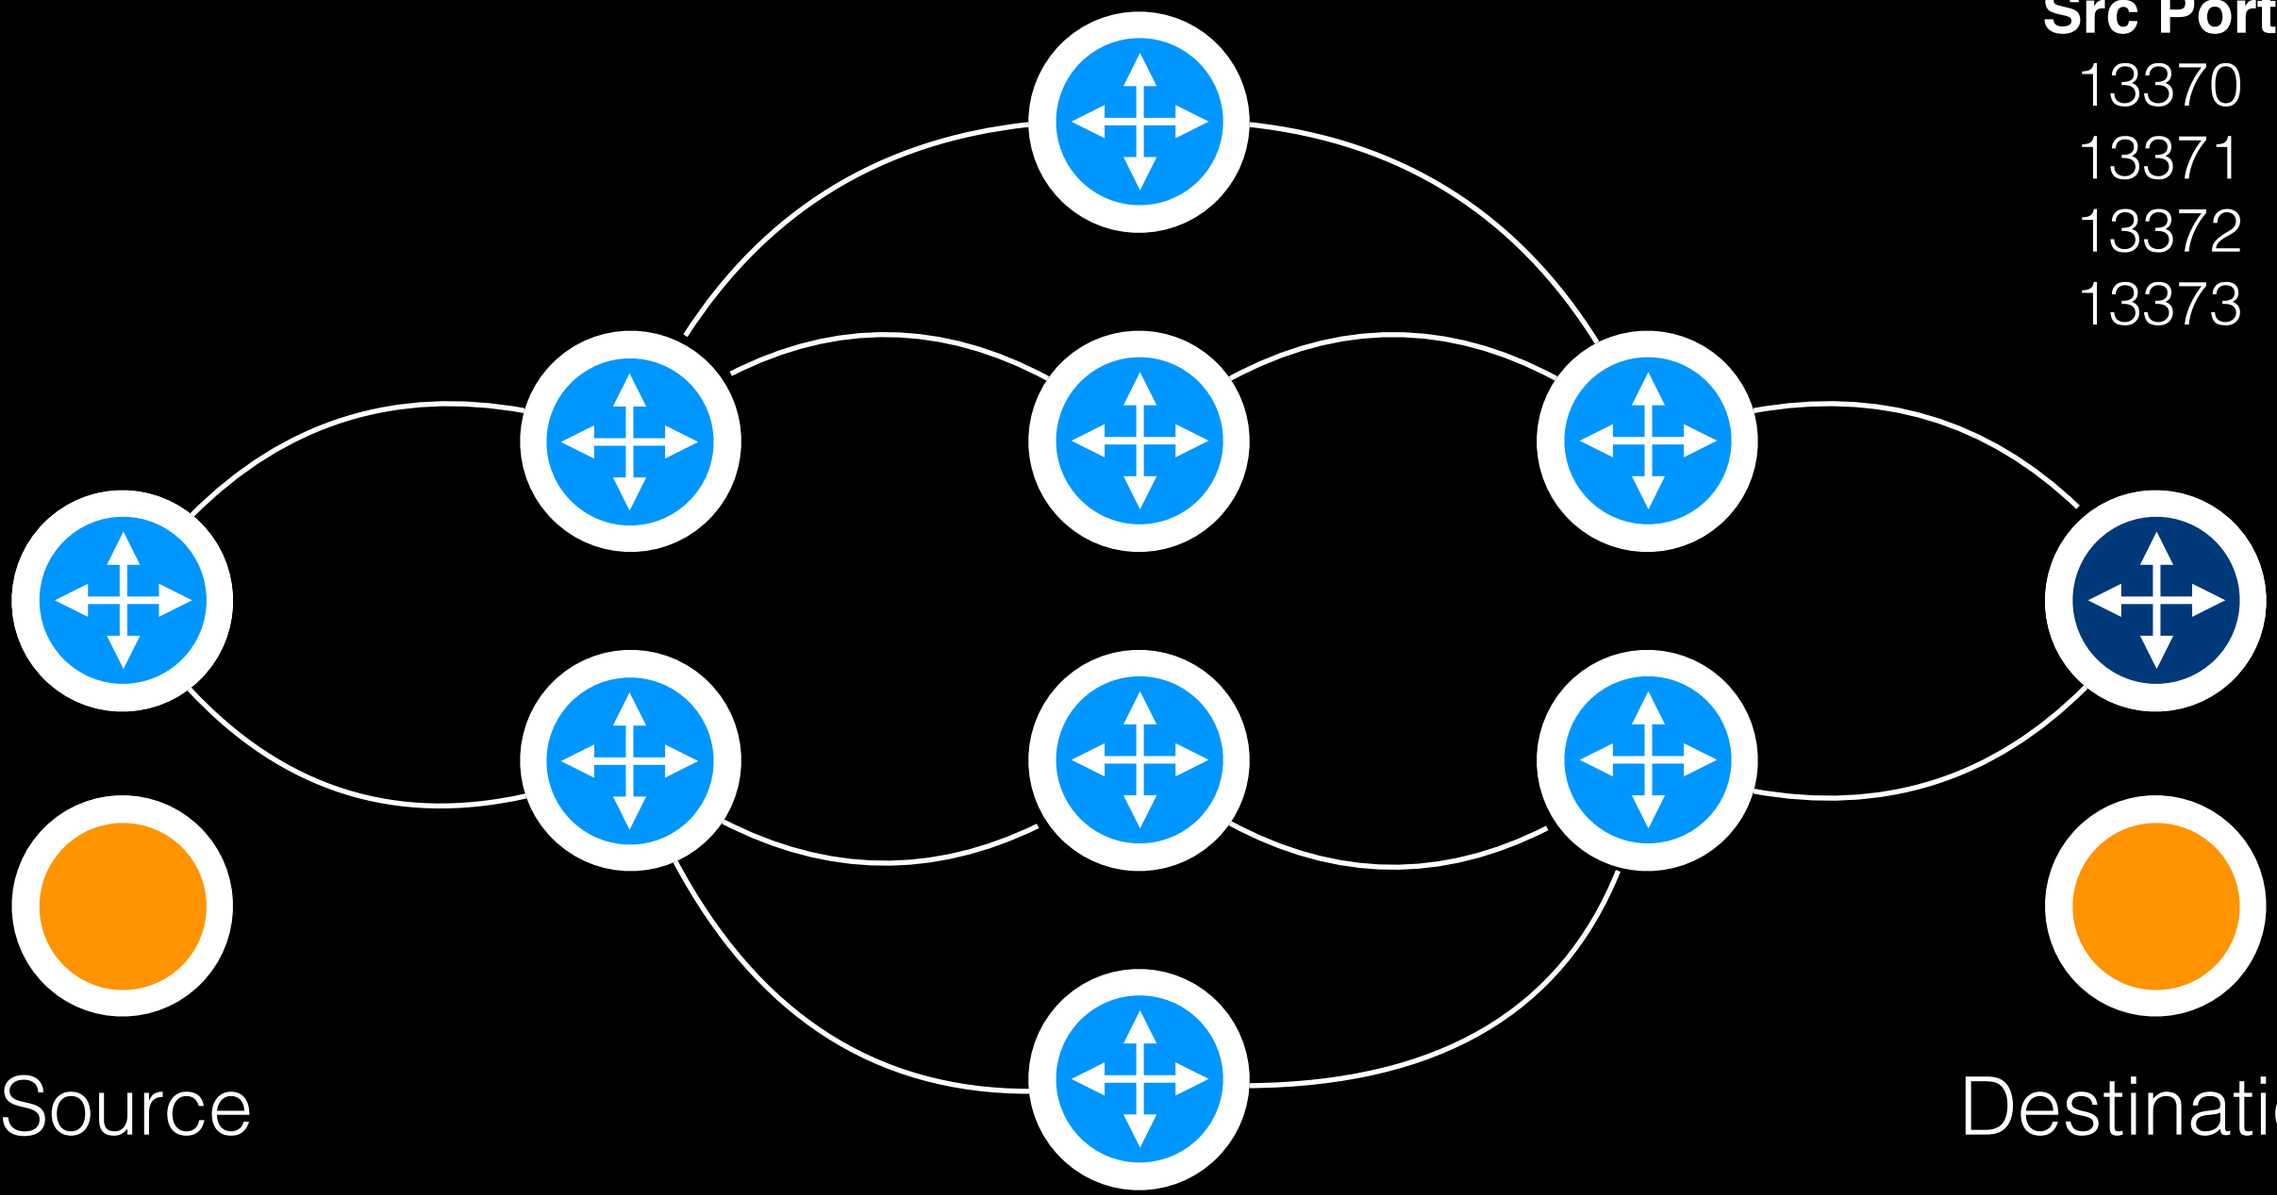

Source

Destination

DC Network Fault Isolation

Isolating Network Faults

Big Fat Fabrics

- Over 1,000 L3 Links between devices in different DCs in the SAME Region
- We know there is loss between hosts, but where?

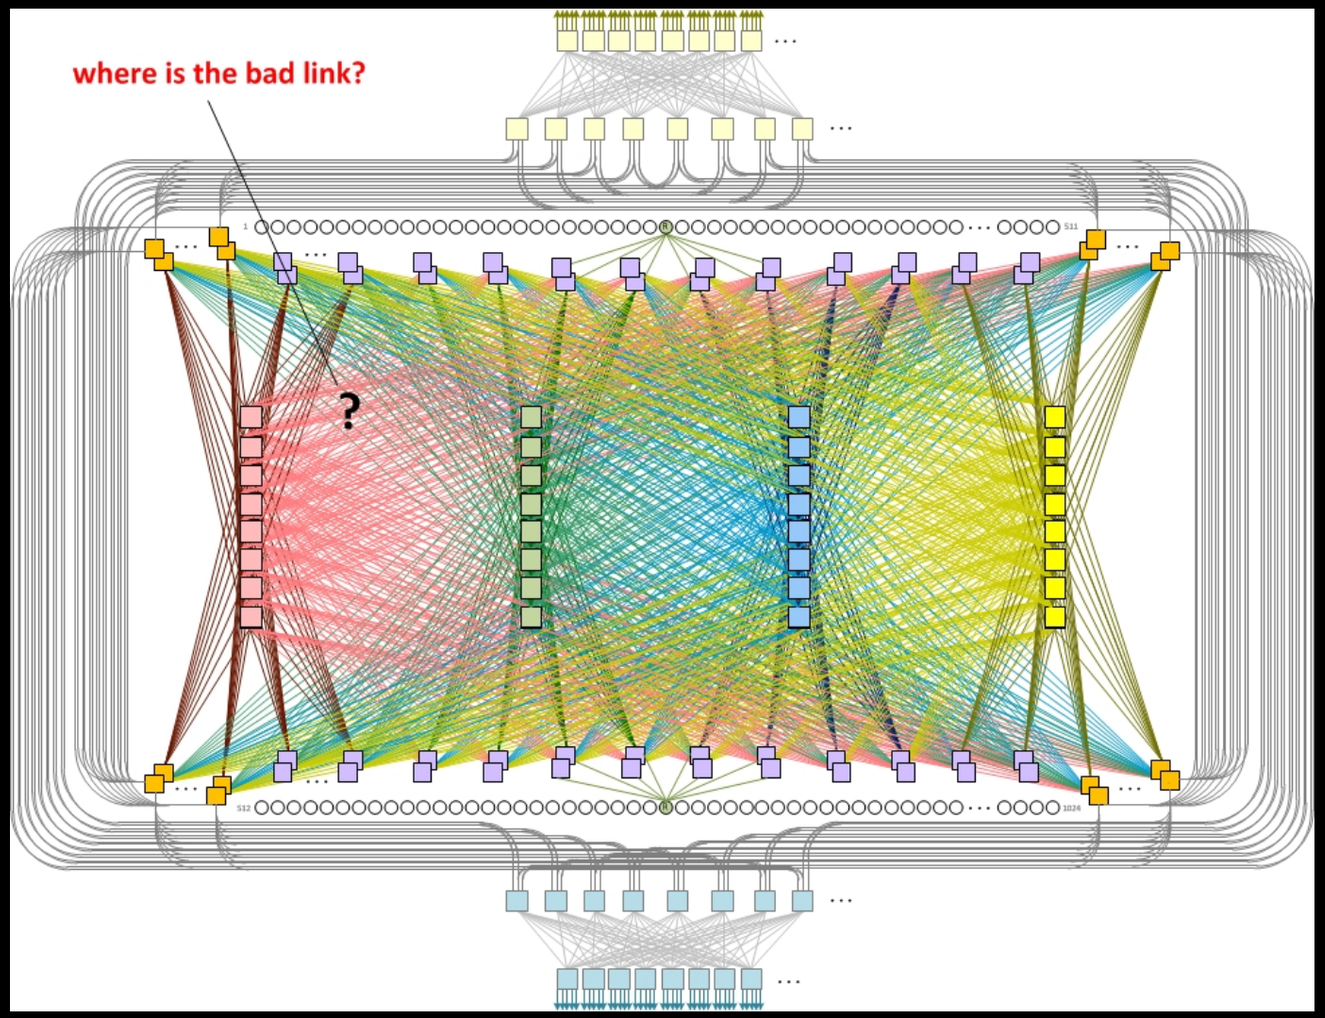

DC Network Fault Isolation

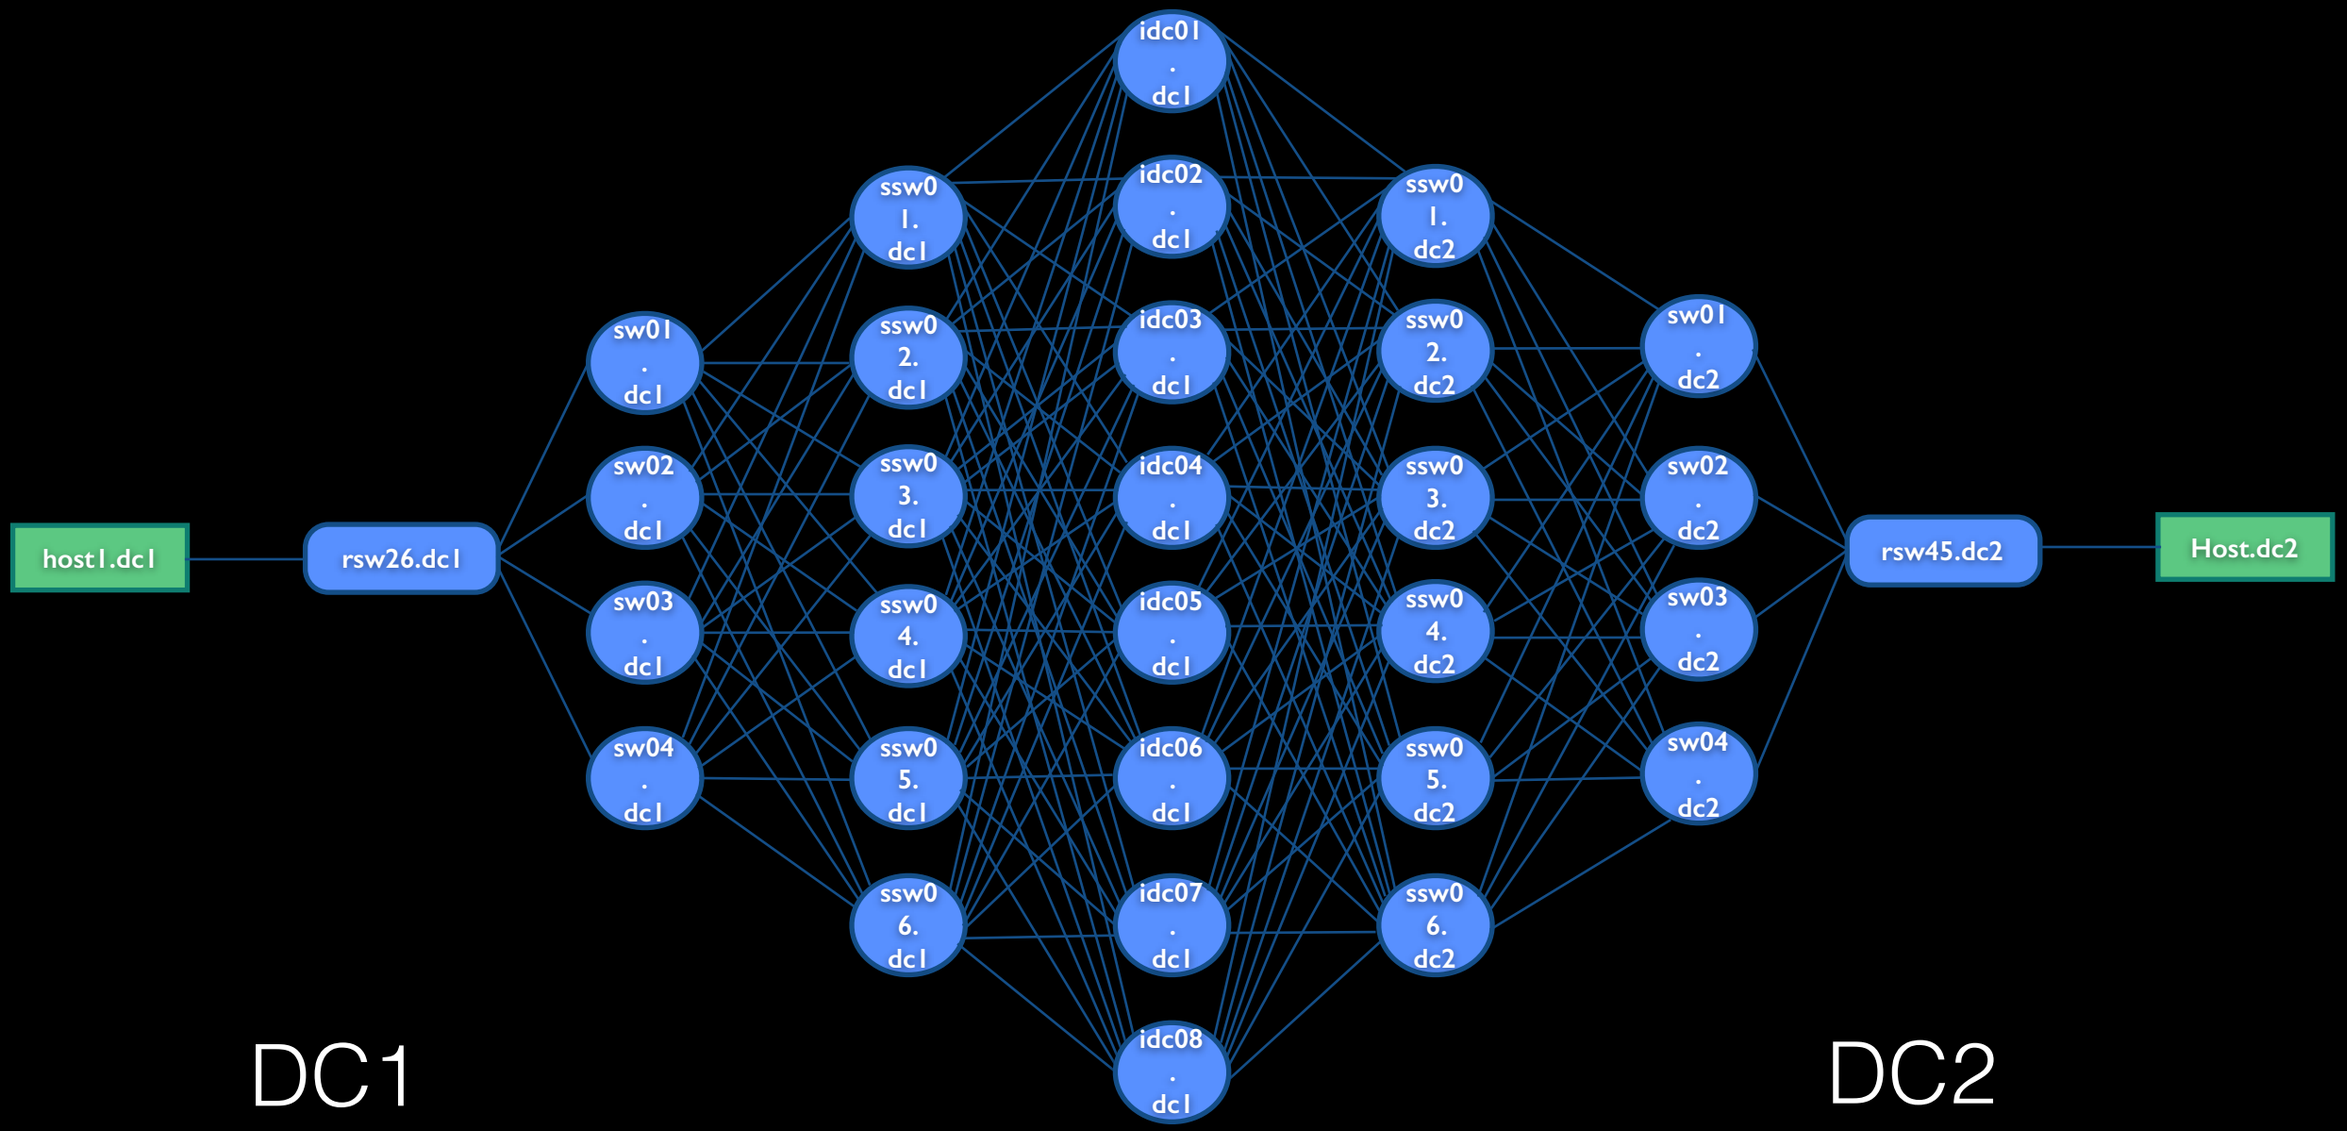

Bad Fabric Card

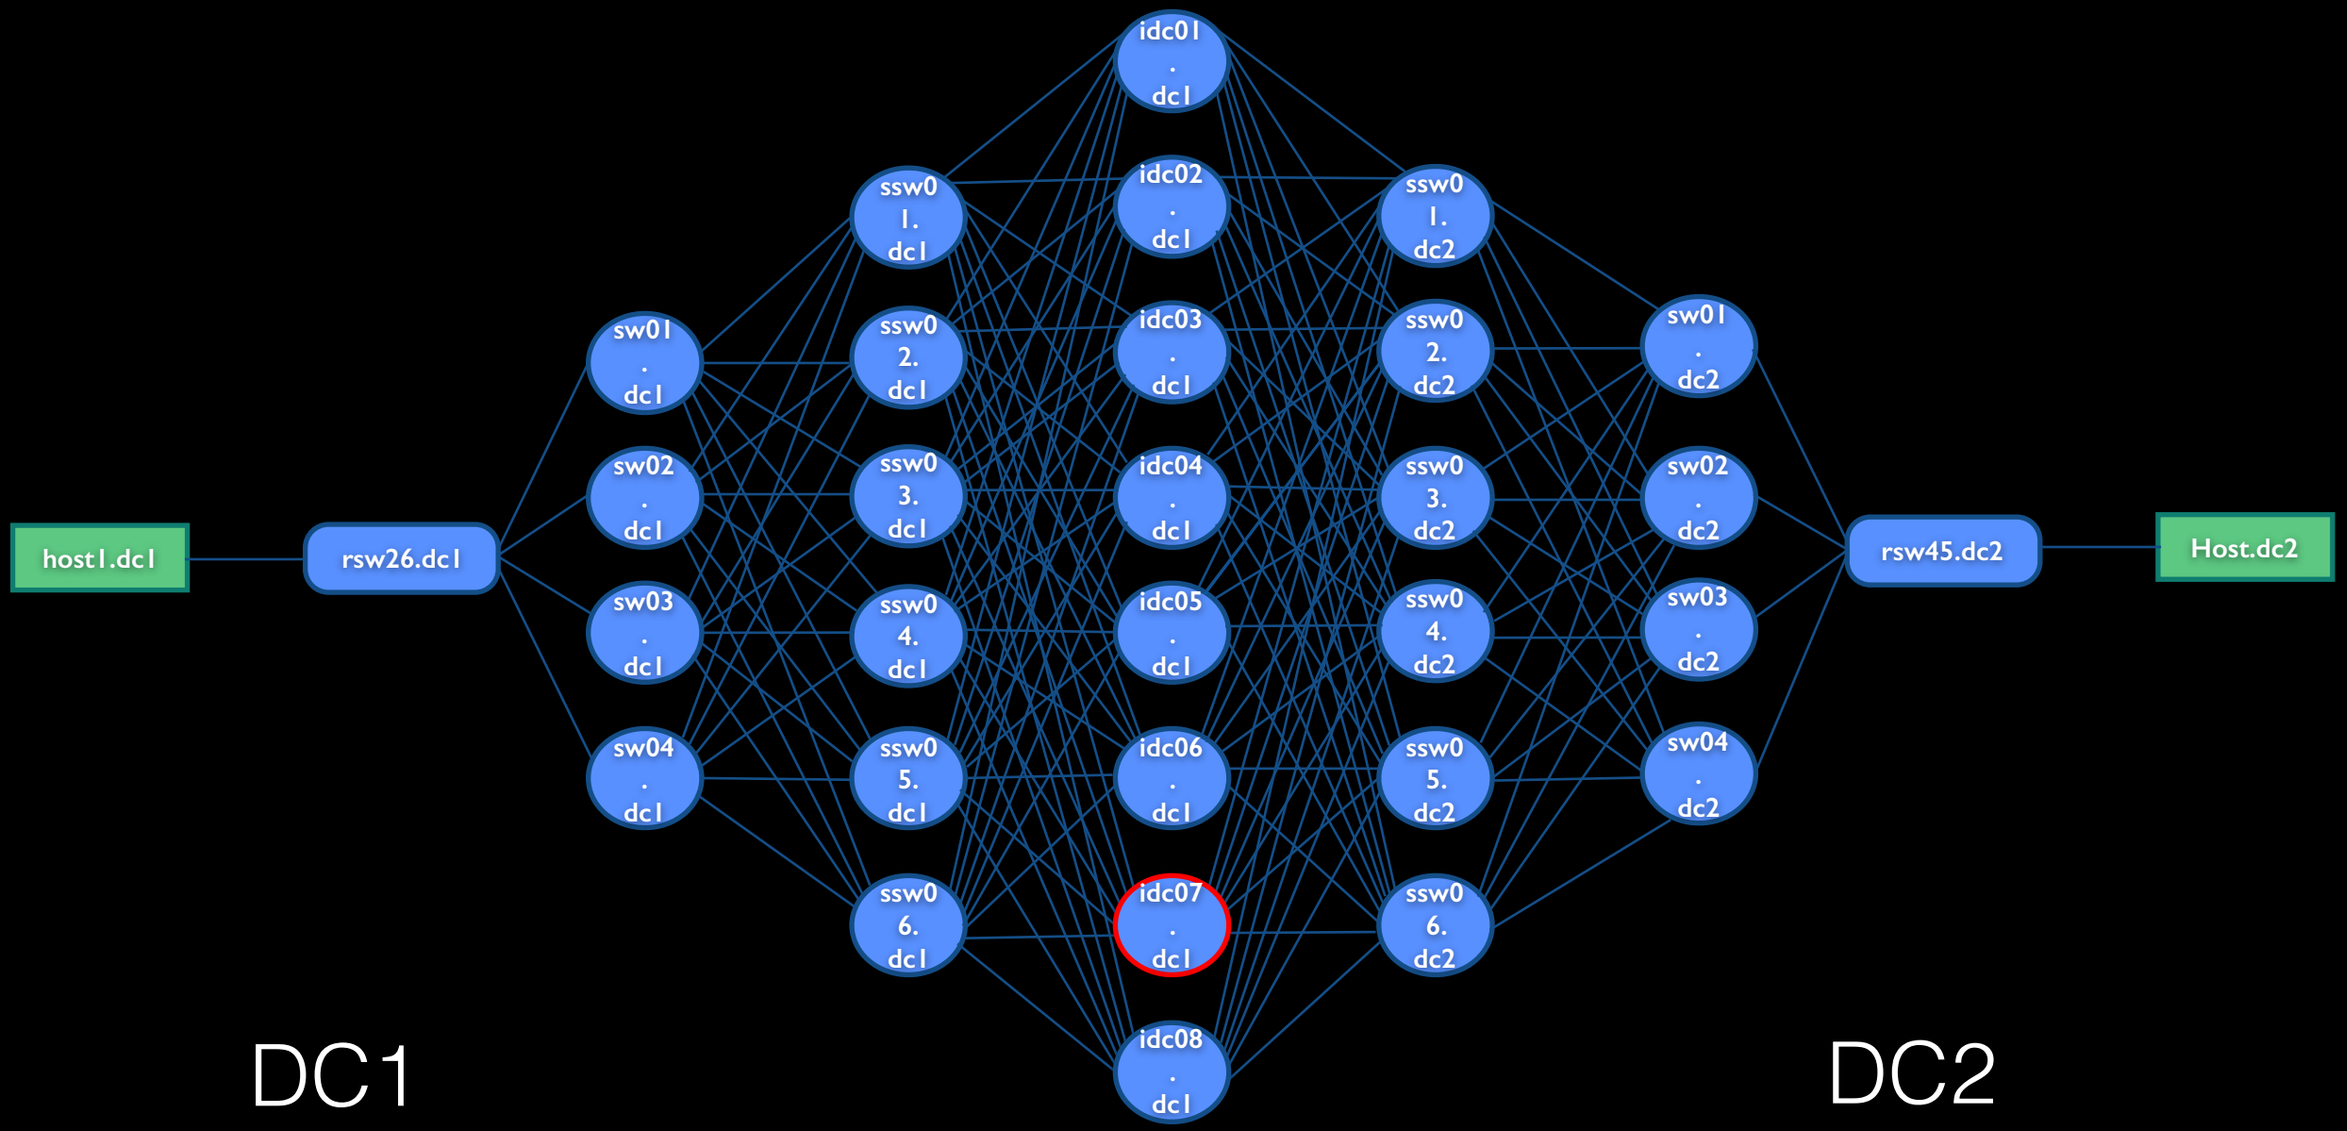

Bad Link

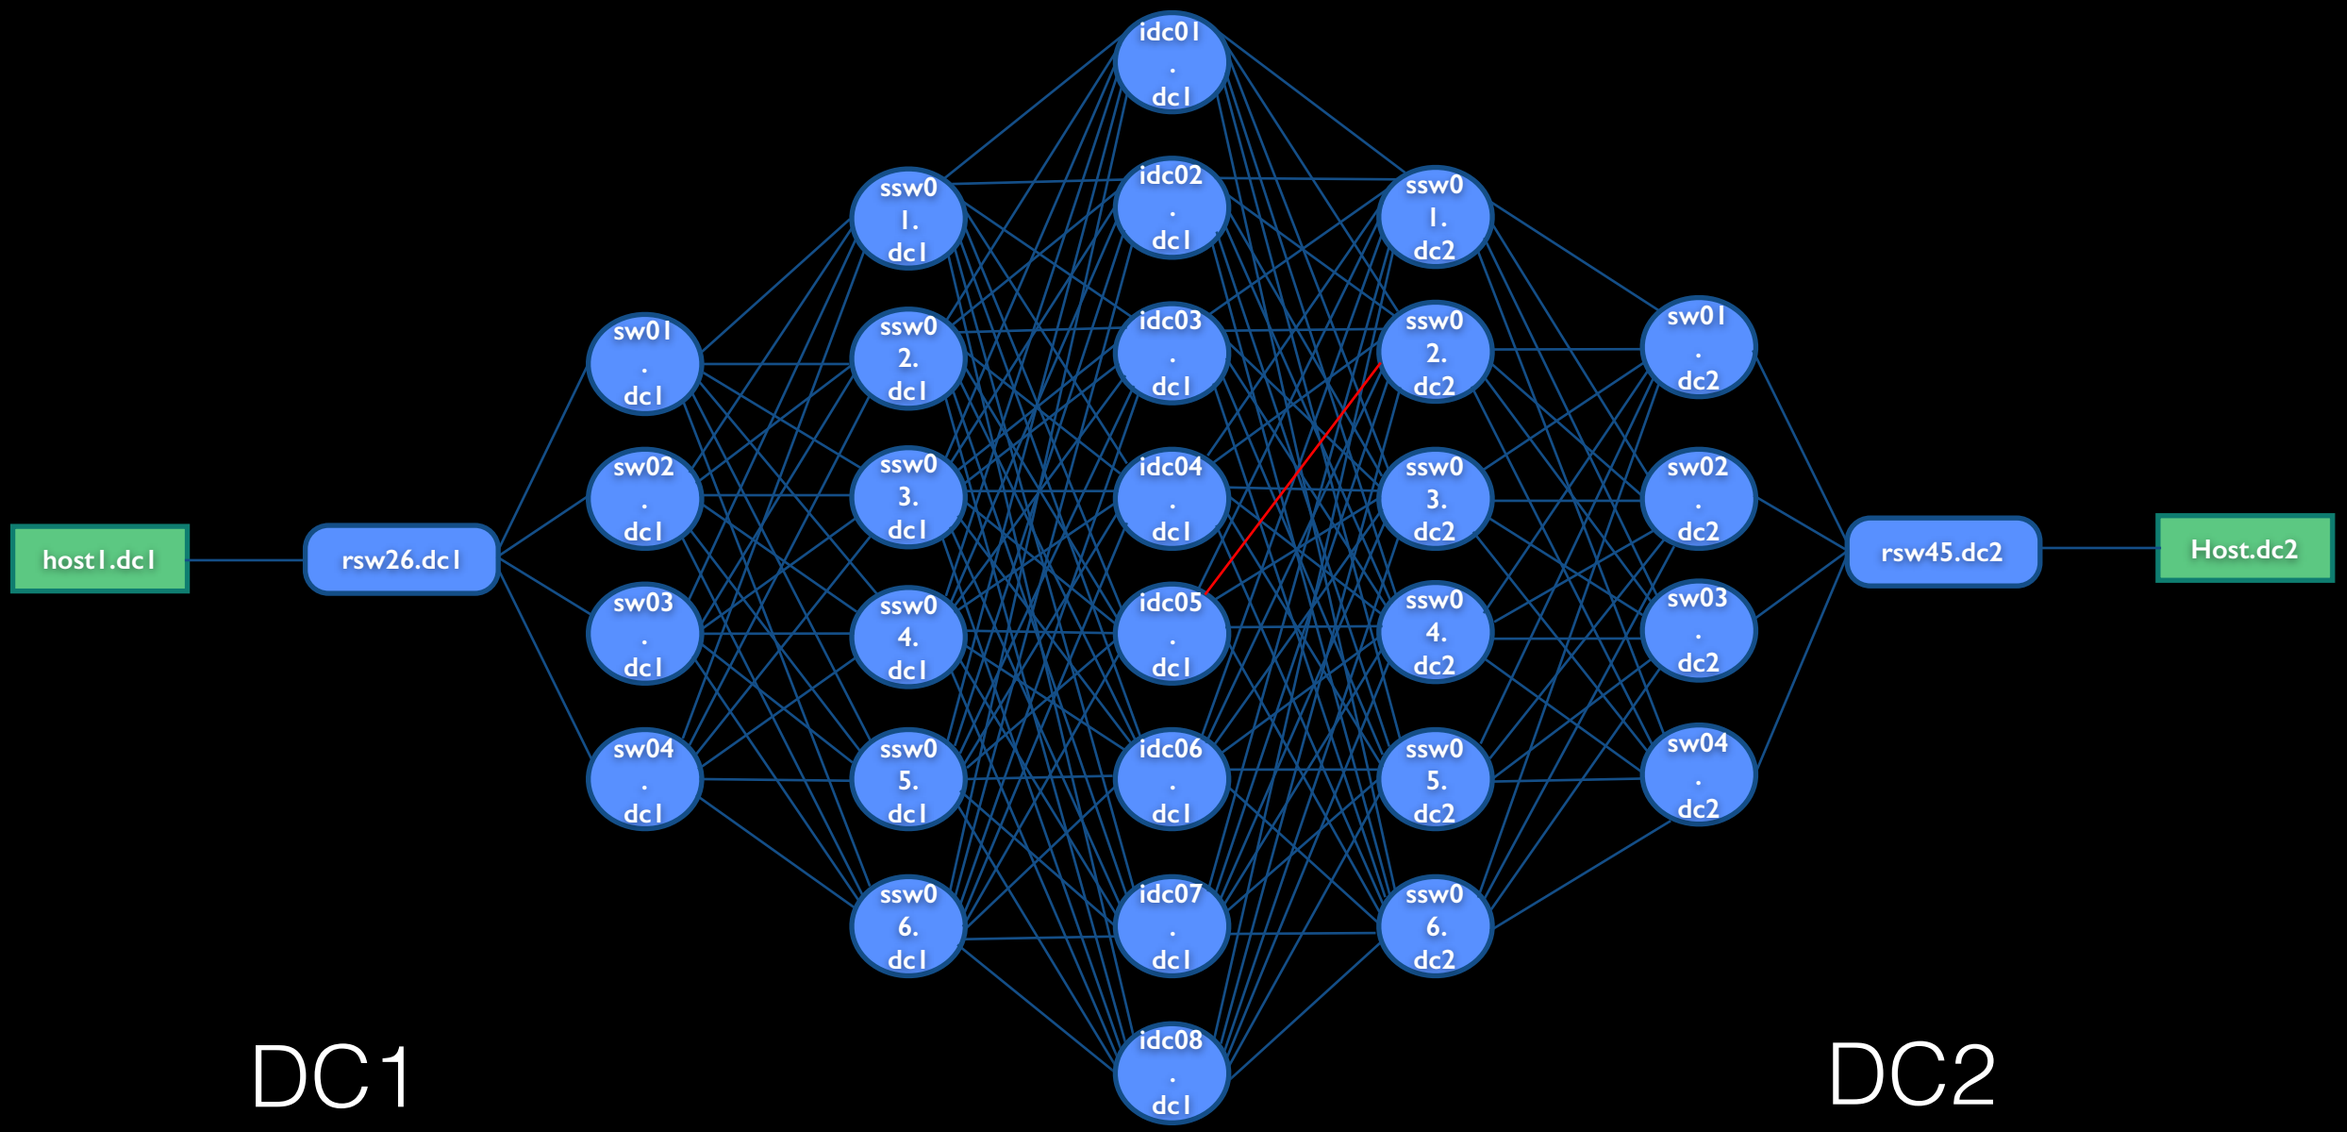
idc01 . dc1
idc02 . dc1
idc03 . dc1
idc04 . dc1
idc05 . dc1
idc06 . dc1
idc07 . dc1
idc08 . dc1
ssw0 1. dc1
ssw0 2. dc1
ssw0 3. dc1
ssw0 4. dc1
ssw0 5. dc1
ssw0 6. dc1
sw01 . dc1
sw02 . dc1
sw03 . dc1
sw04 . dc1
host1.dc1
rsw26.dc1
ssw0 1. dc2
ssw0 2. dc2
ssw0 3. dc2
ssw0 4. dc2
ssw0 5. dc2
ssw0 6. dc2
sw01 . dc2
sw02 . dc2
sw03 . dc2
sw04 . dc2
rsw45.dc2
Host.dc2
DC1
DC2

ECMP Packet Hashing

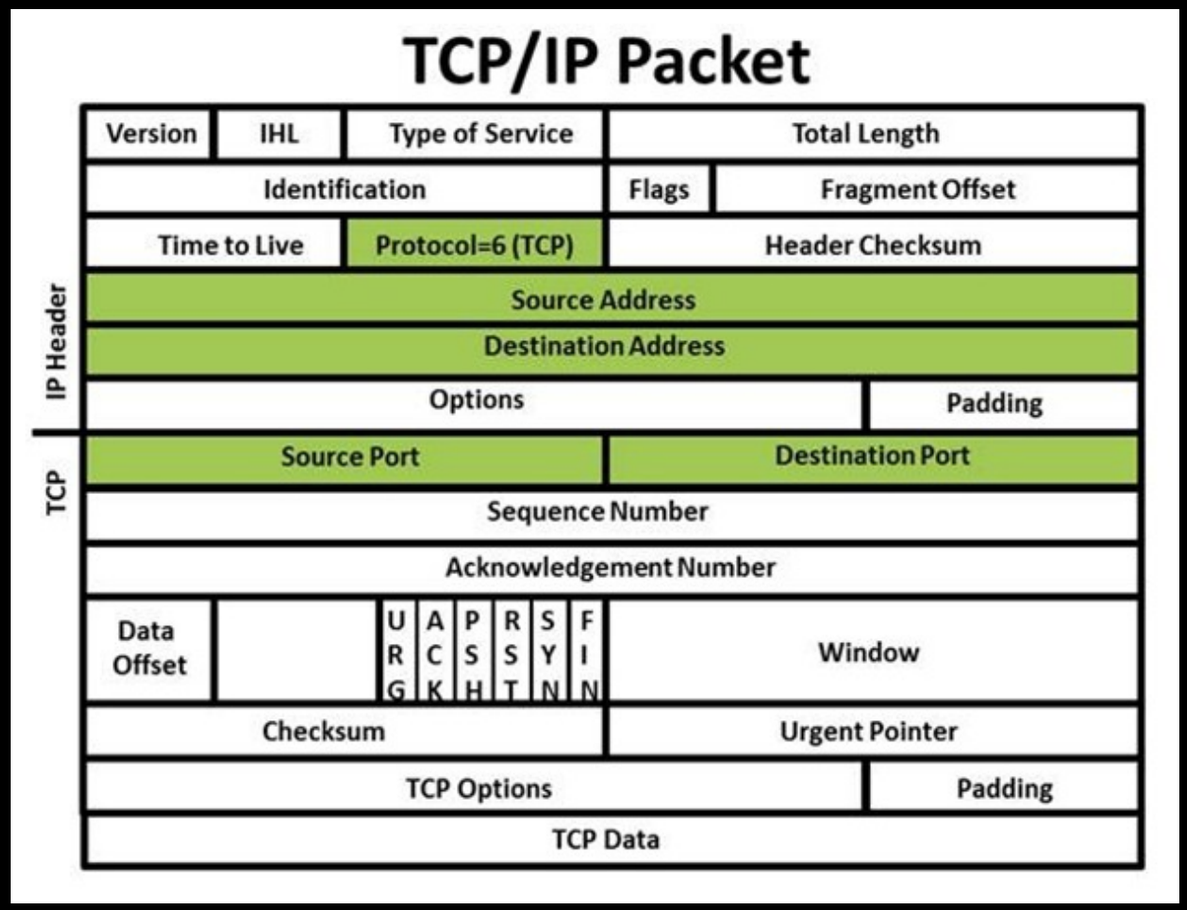

DC Network Fault Isolation

Port Range	50,000 - 54,000
Request Rate	Approx 20,000 min

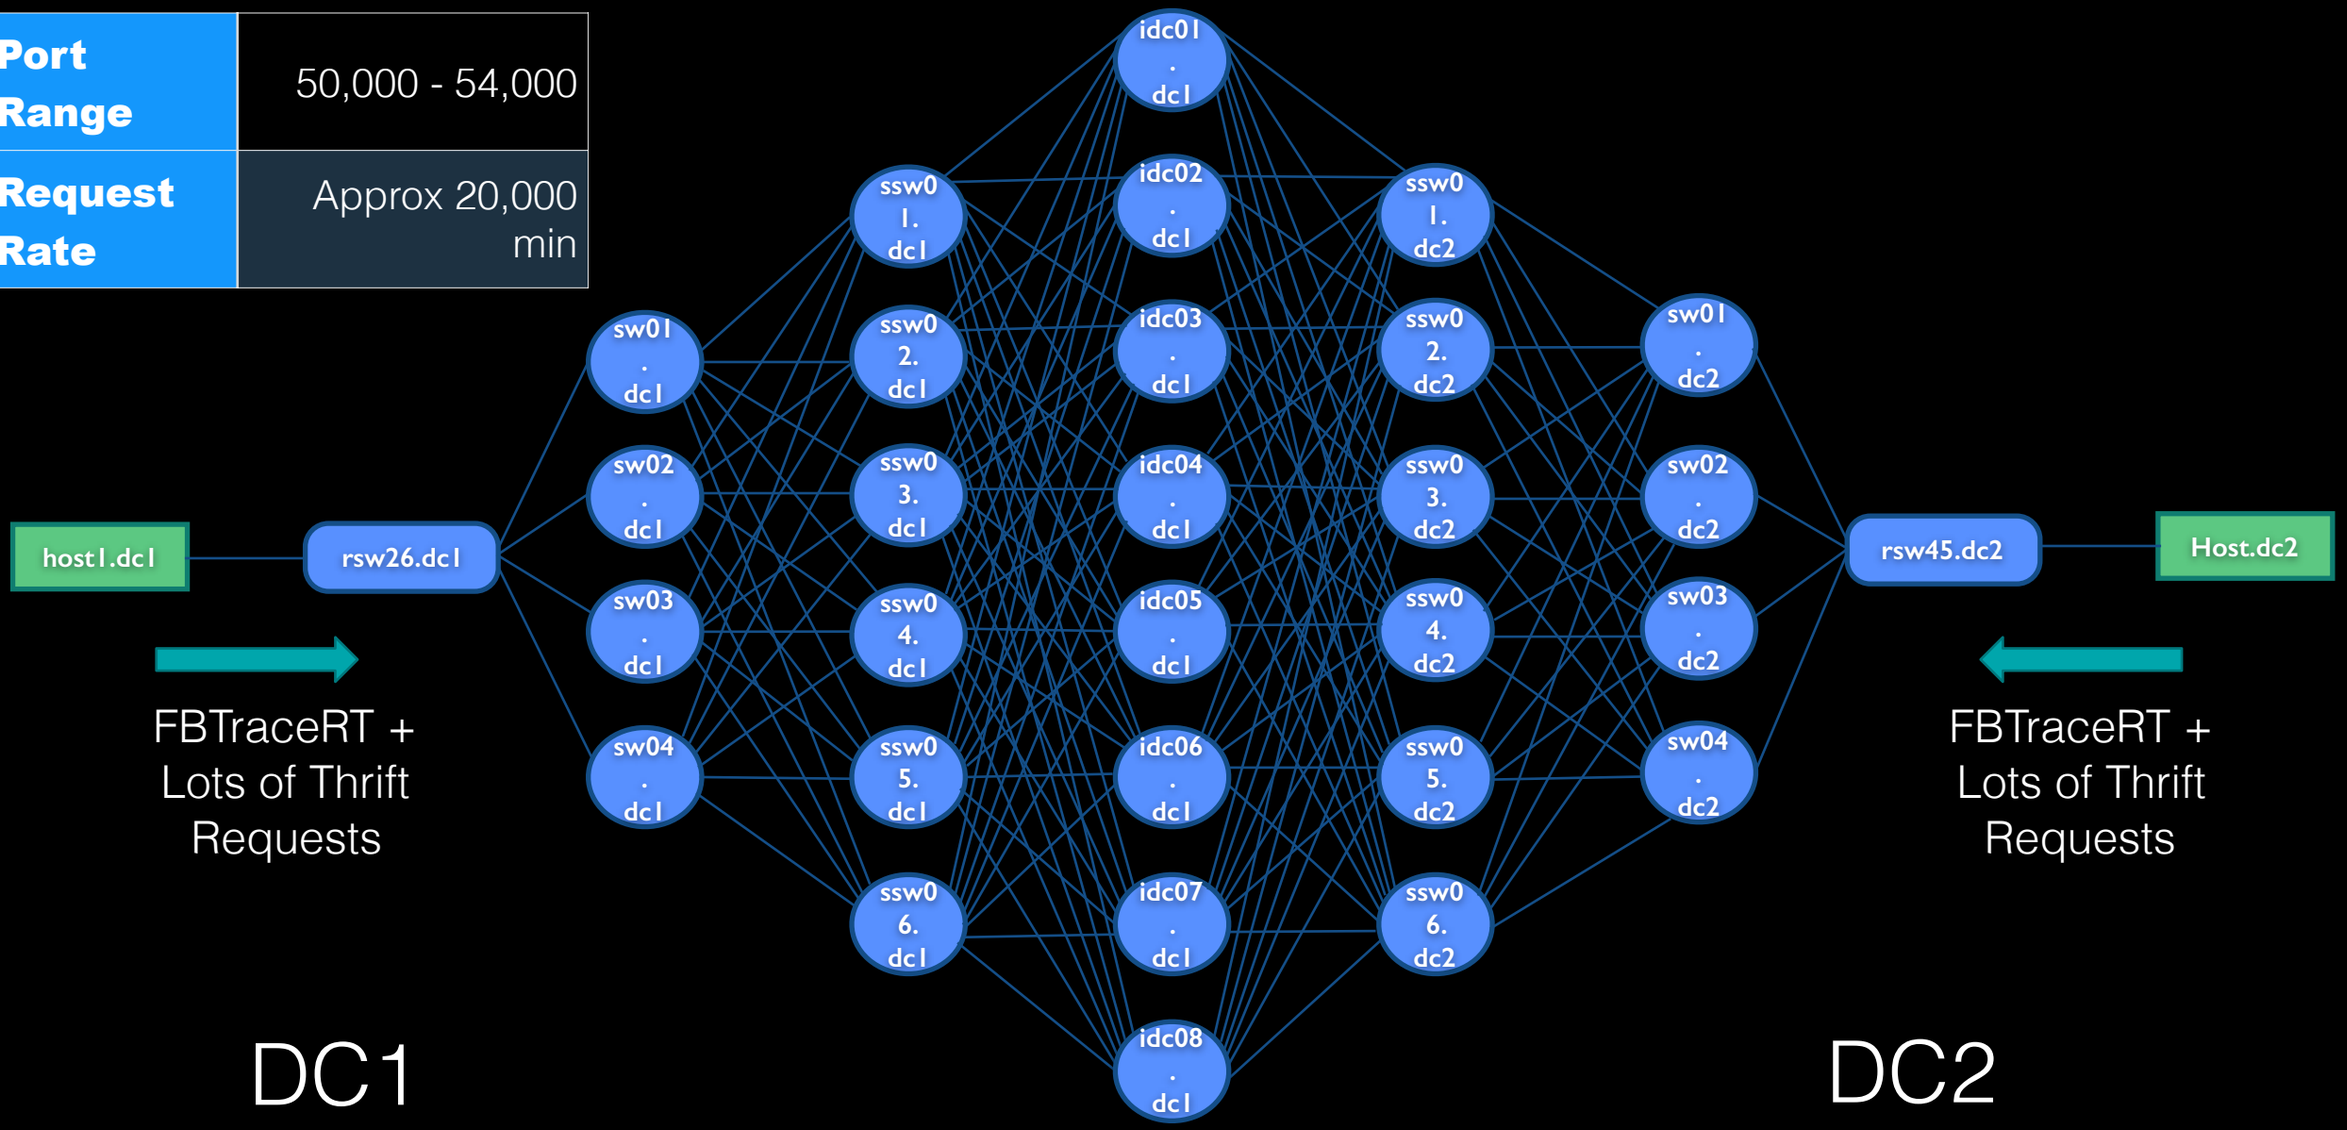

Map The Loss for Each Request

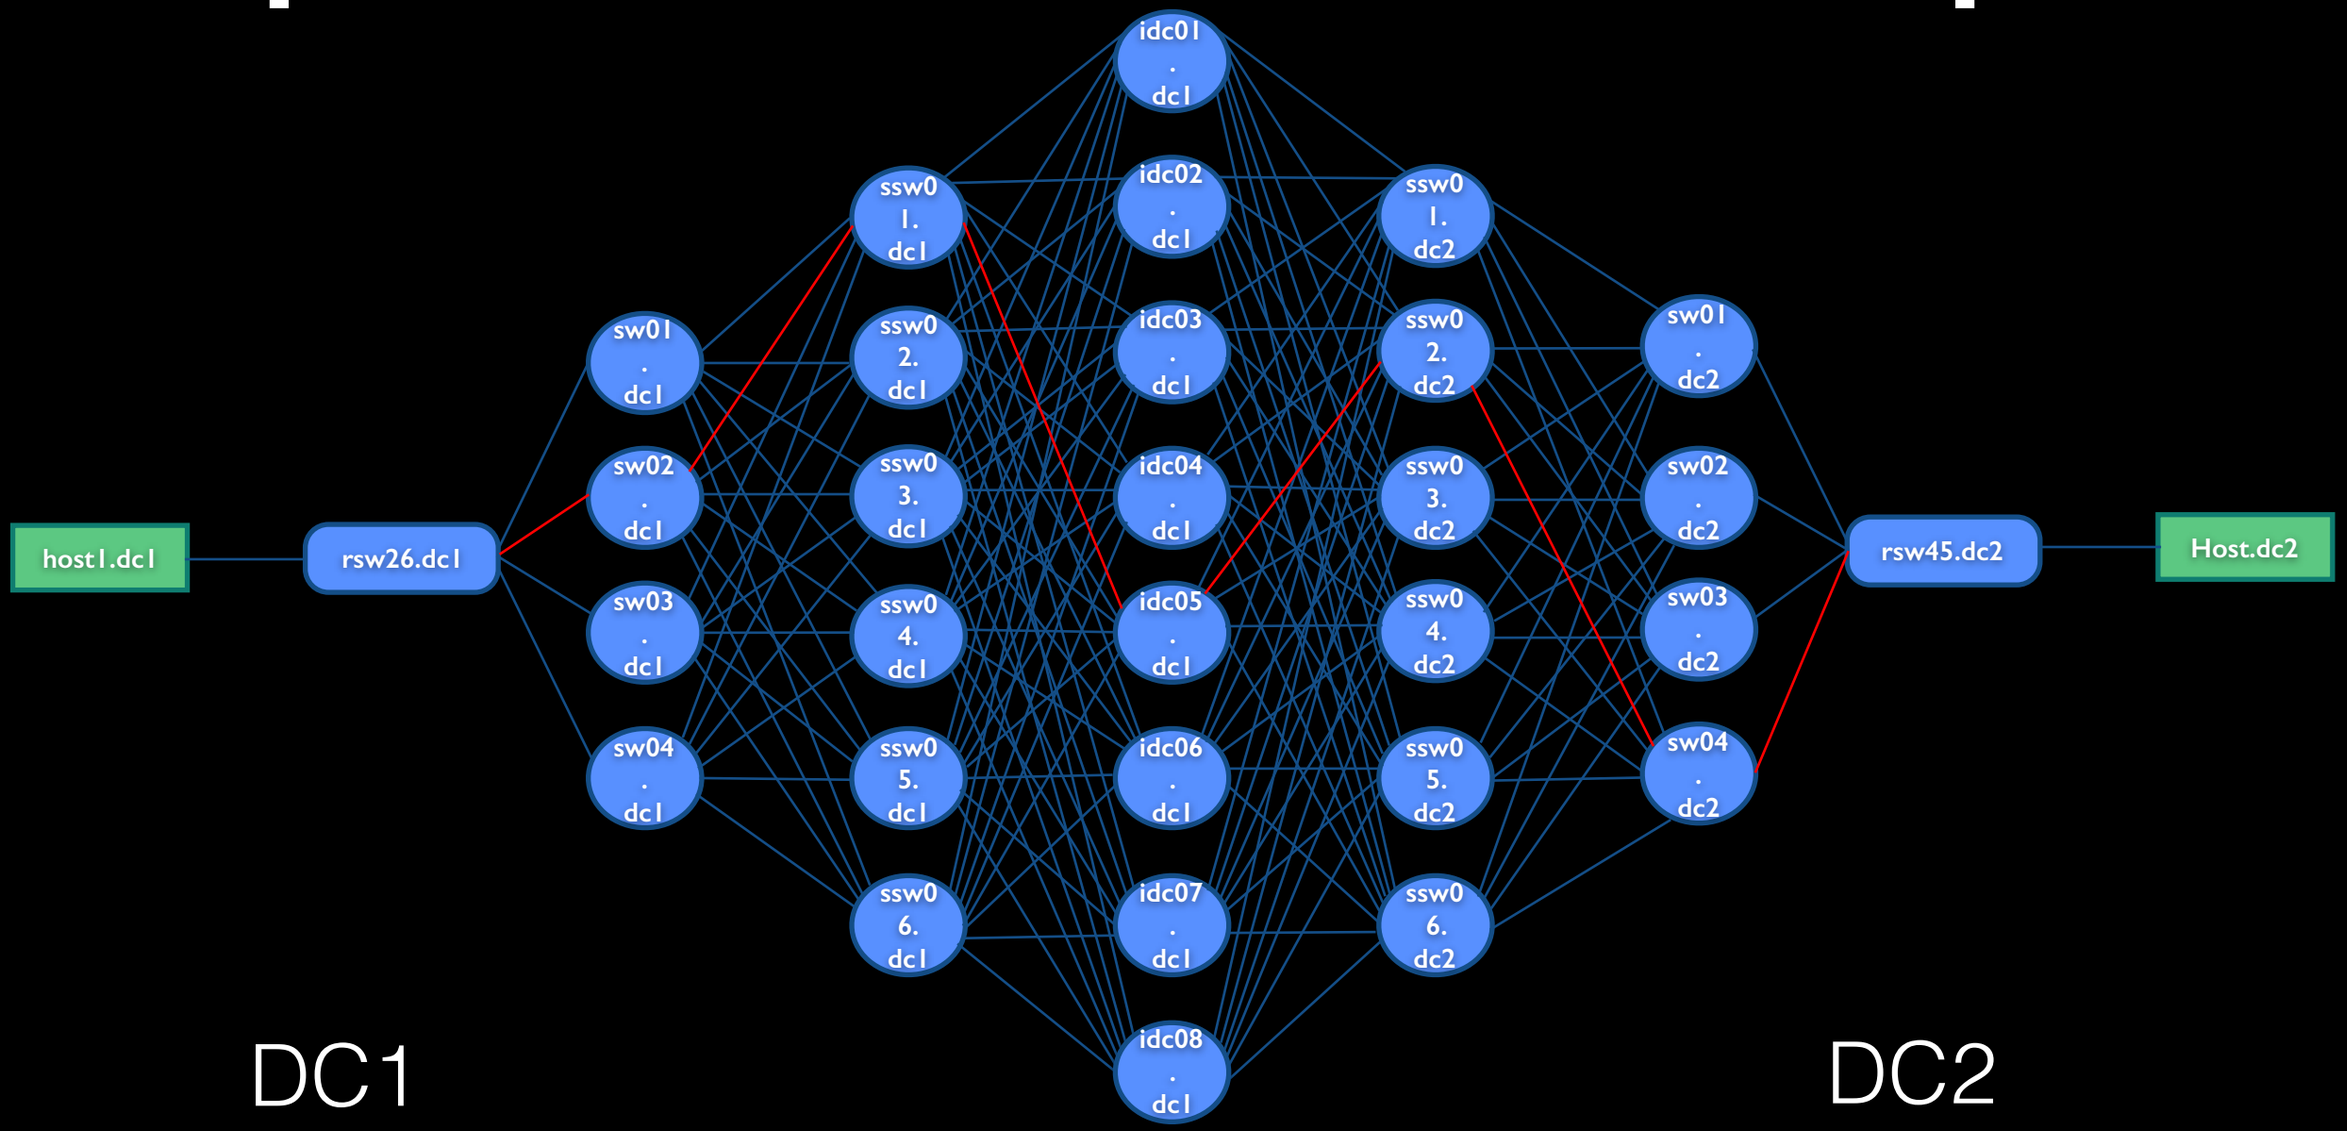

Map The Loss for Each Request

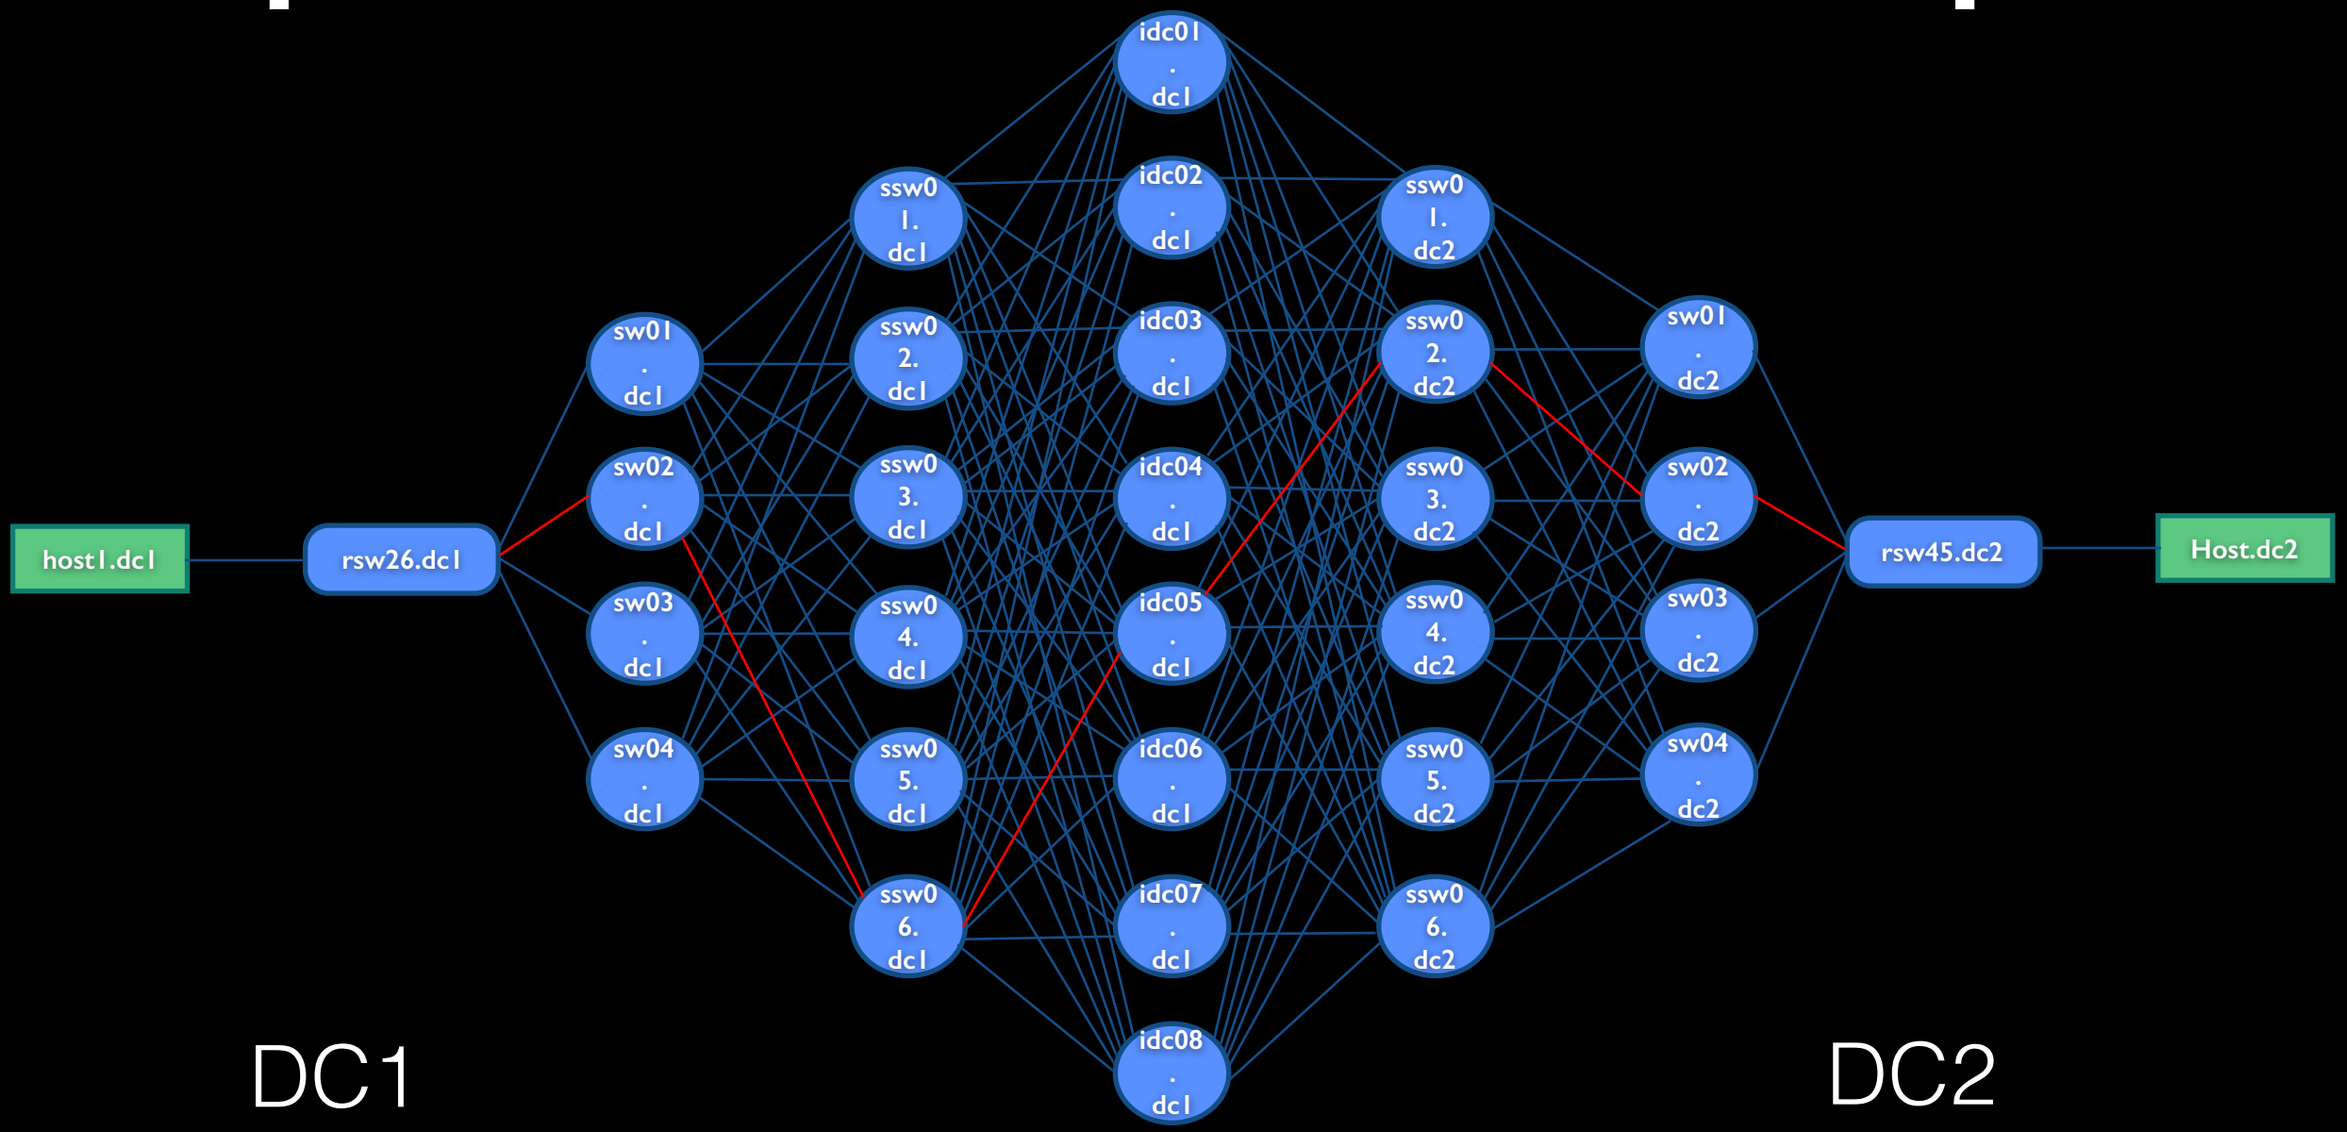

Overlay the Loss on Network Map

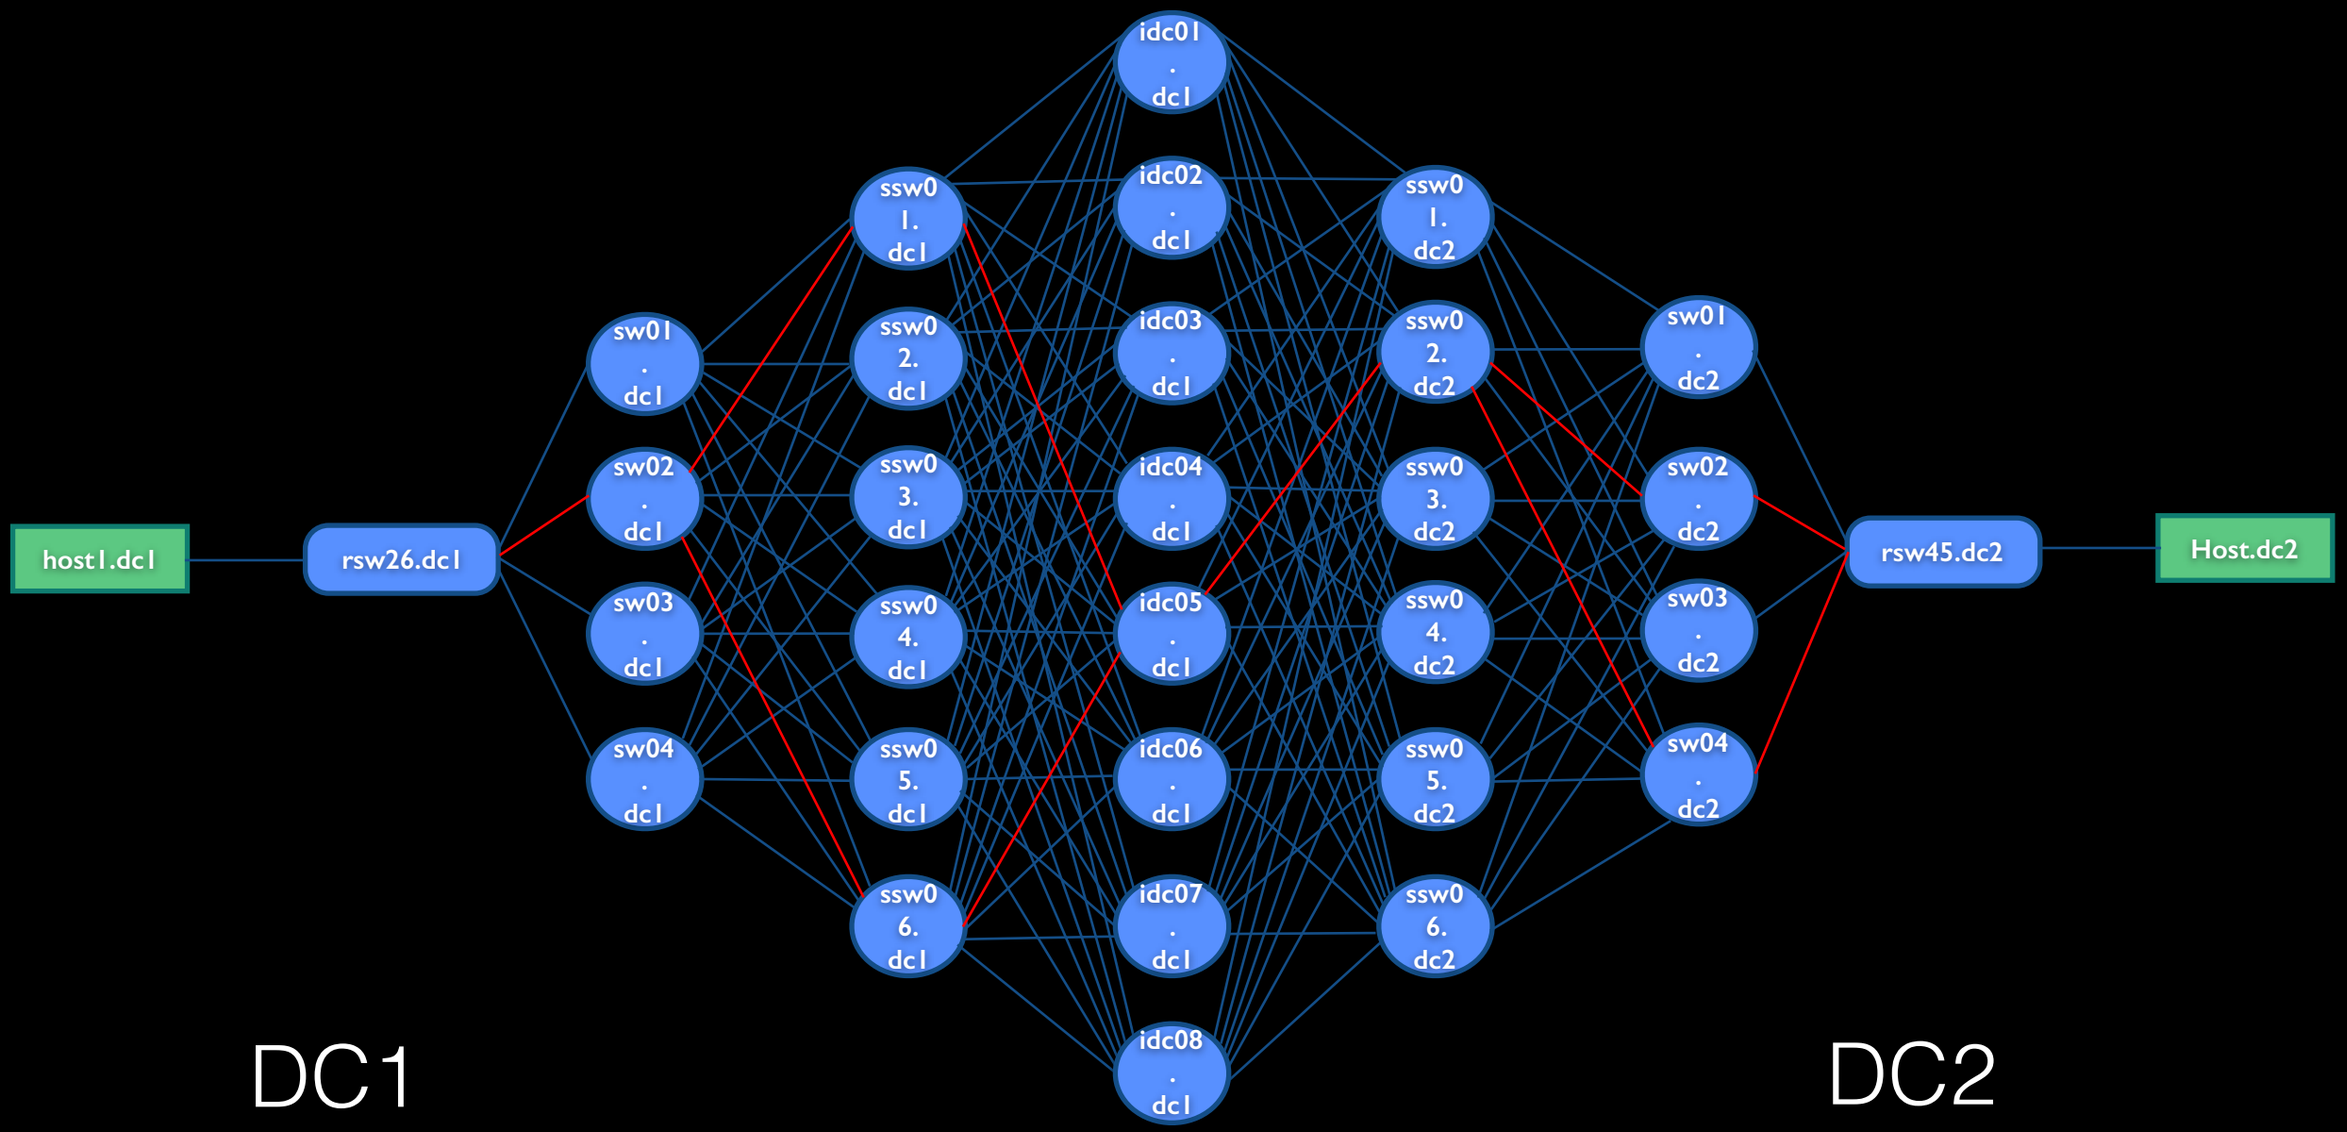

Analyze the Loss

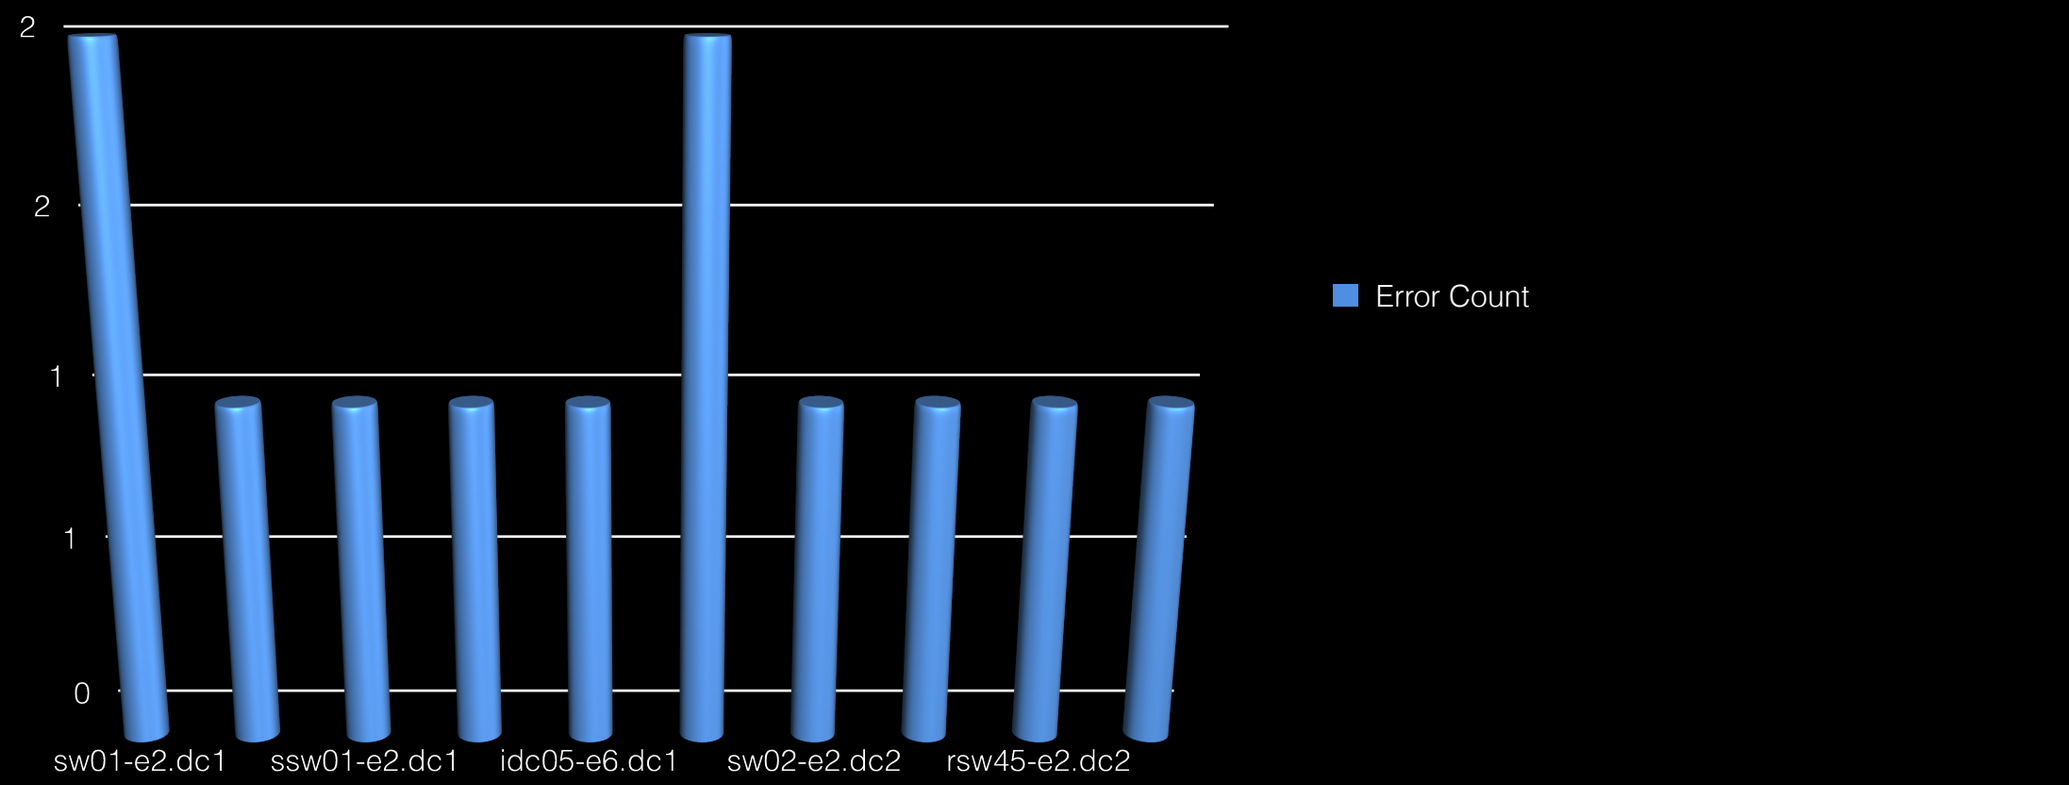

Network Symmetry

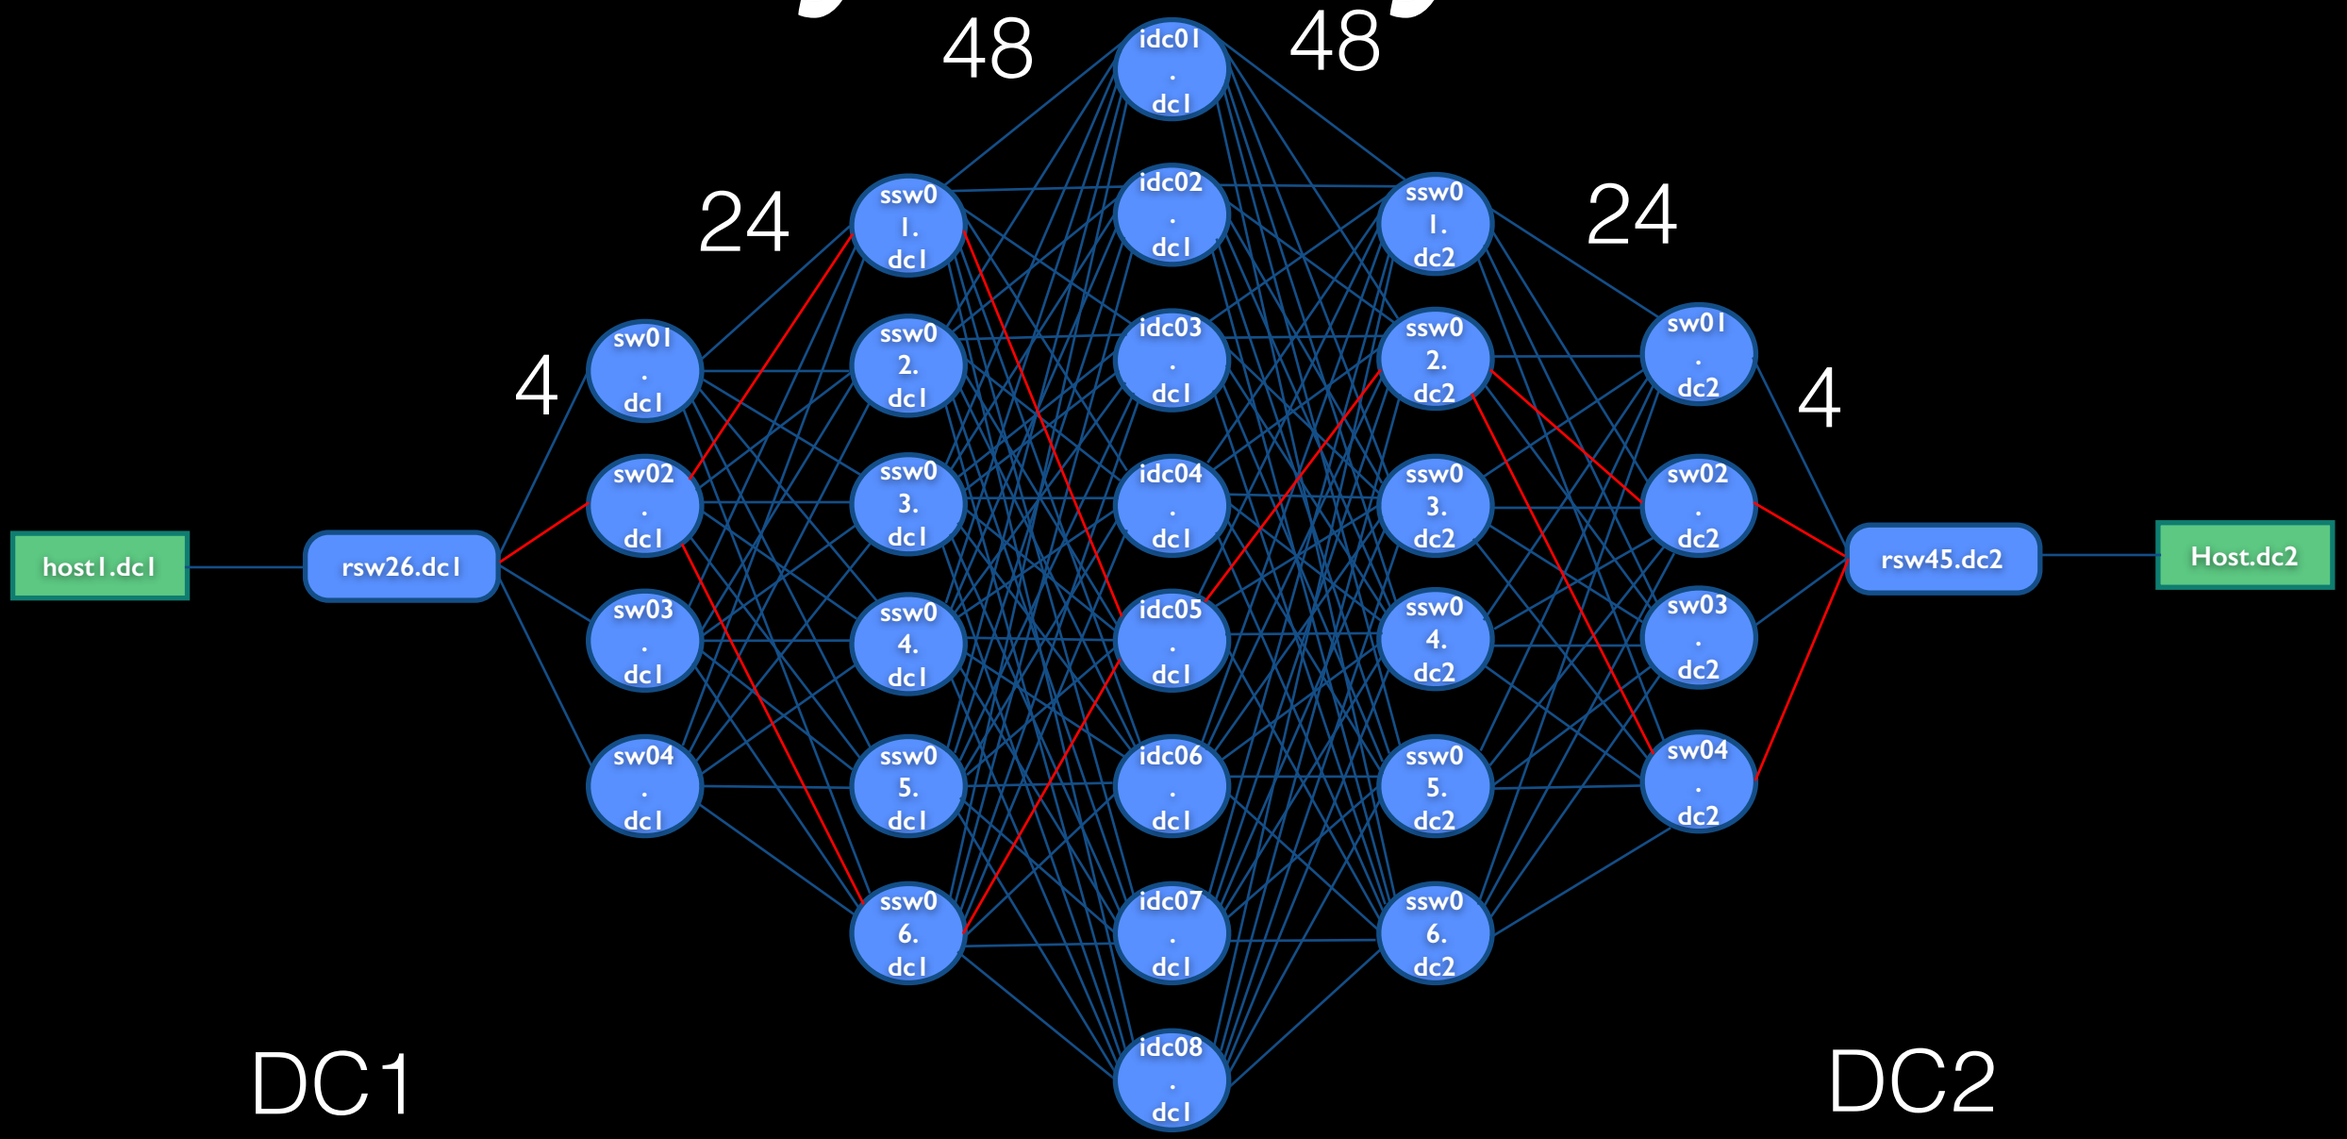

Clearer Signal

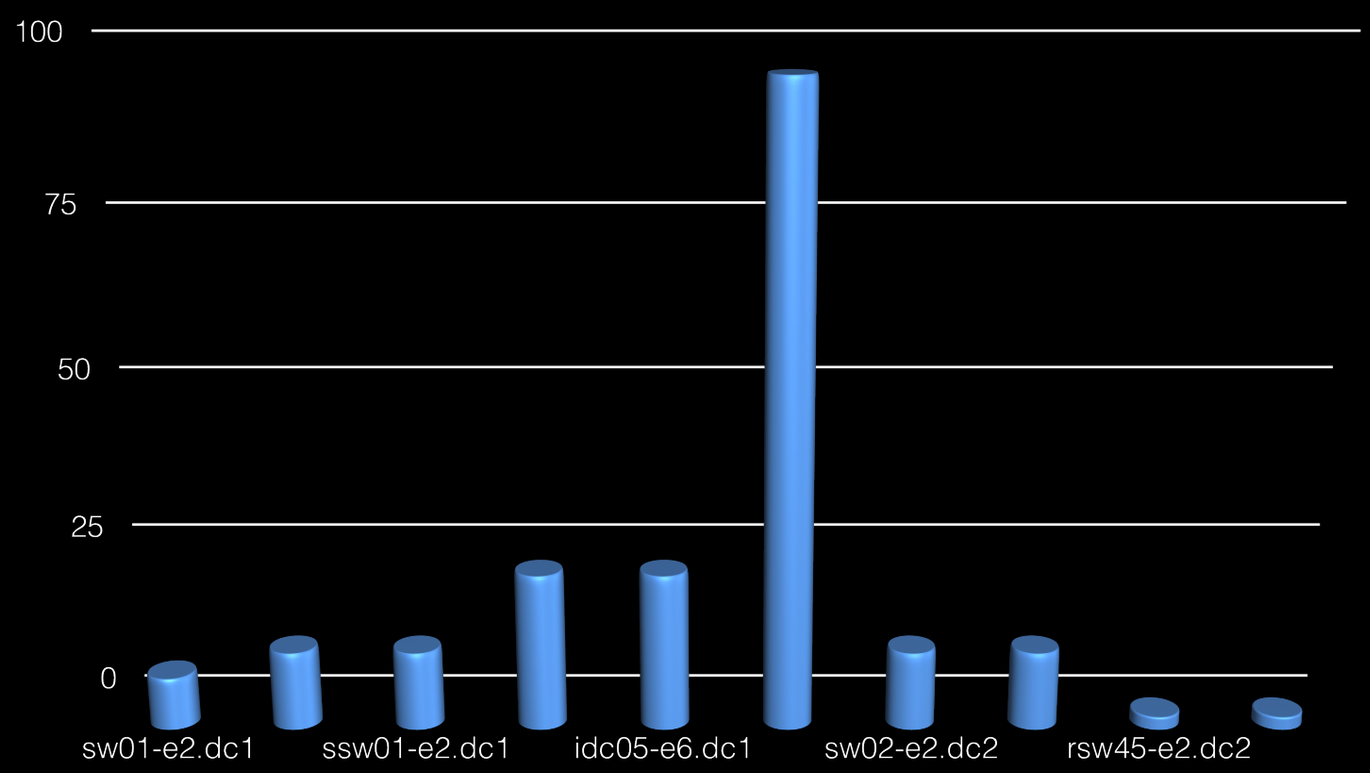

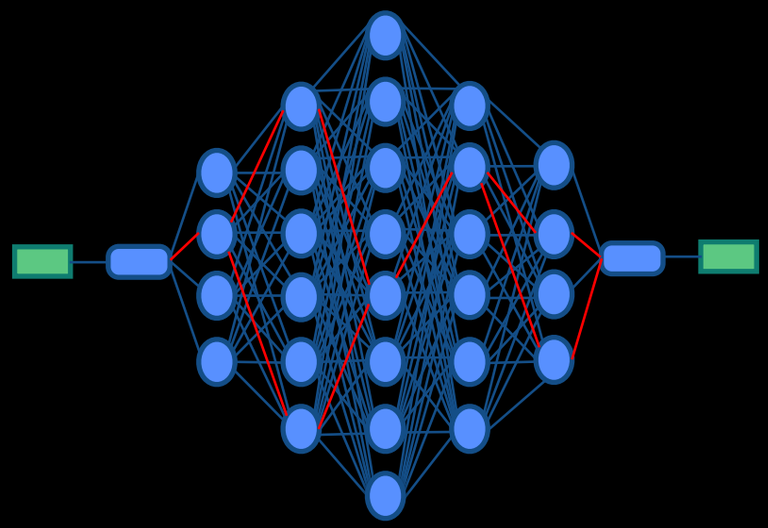

Fabric Grey Failure Detection

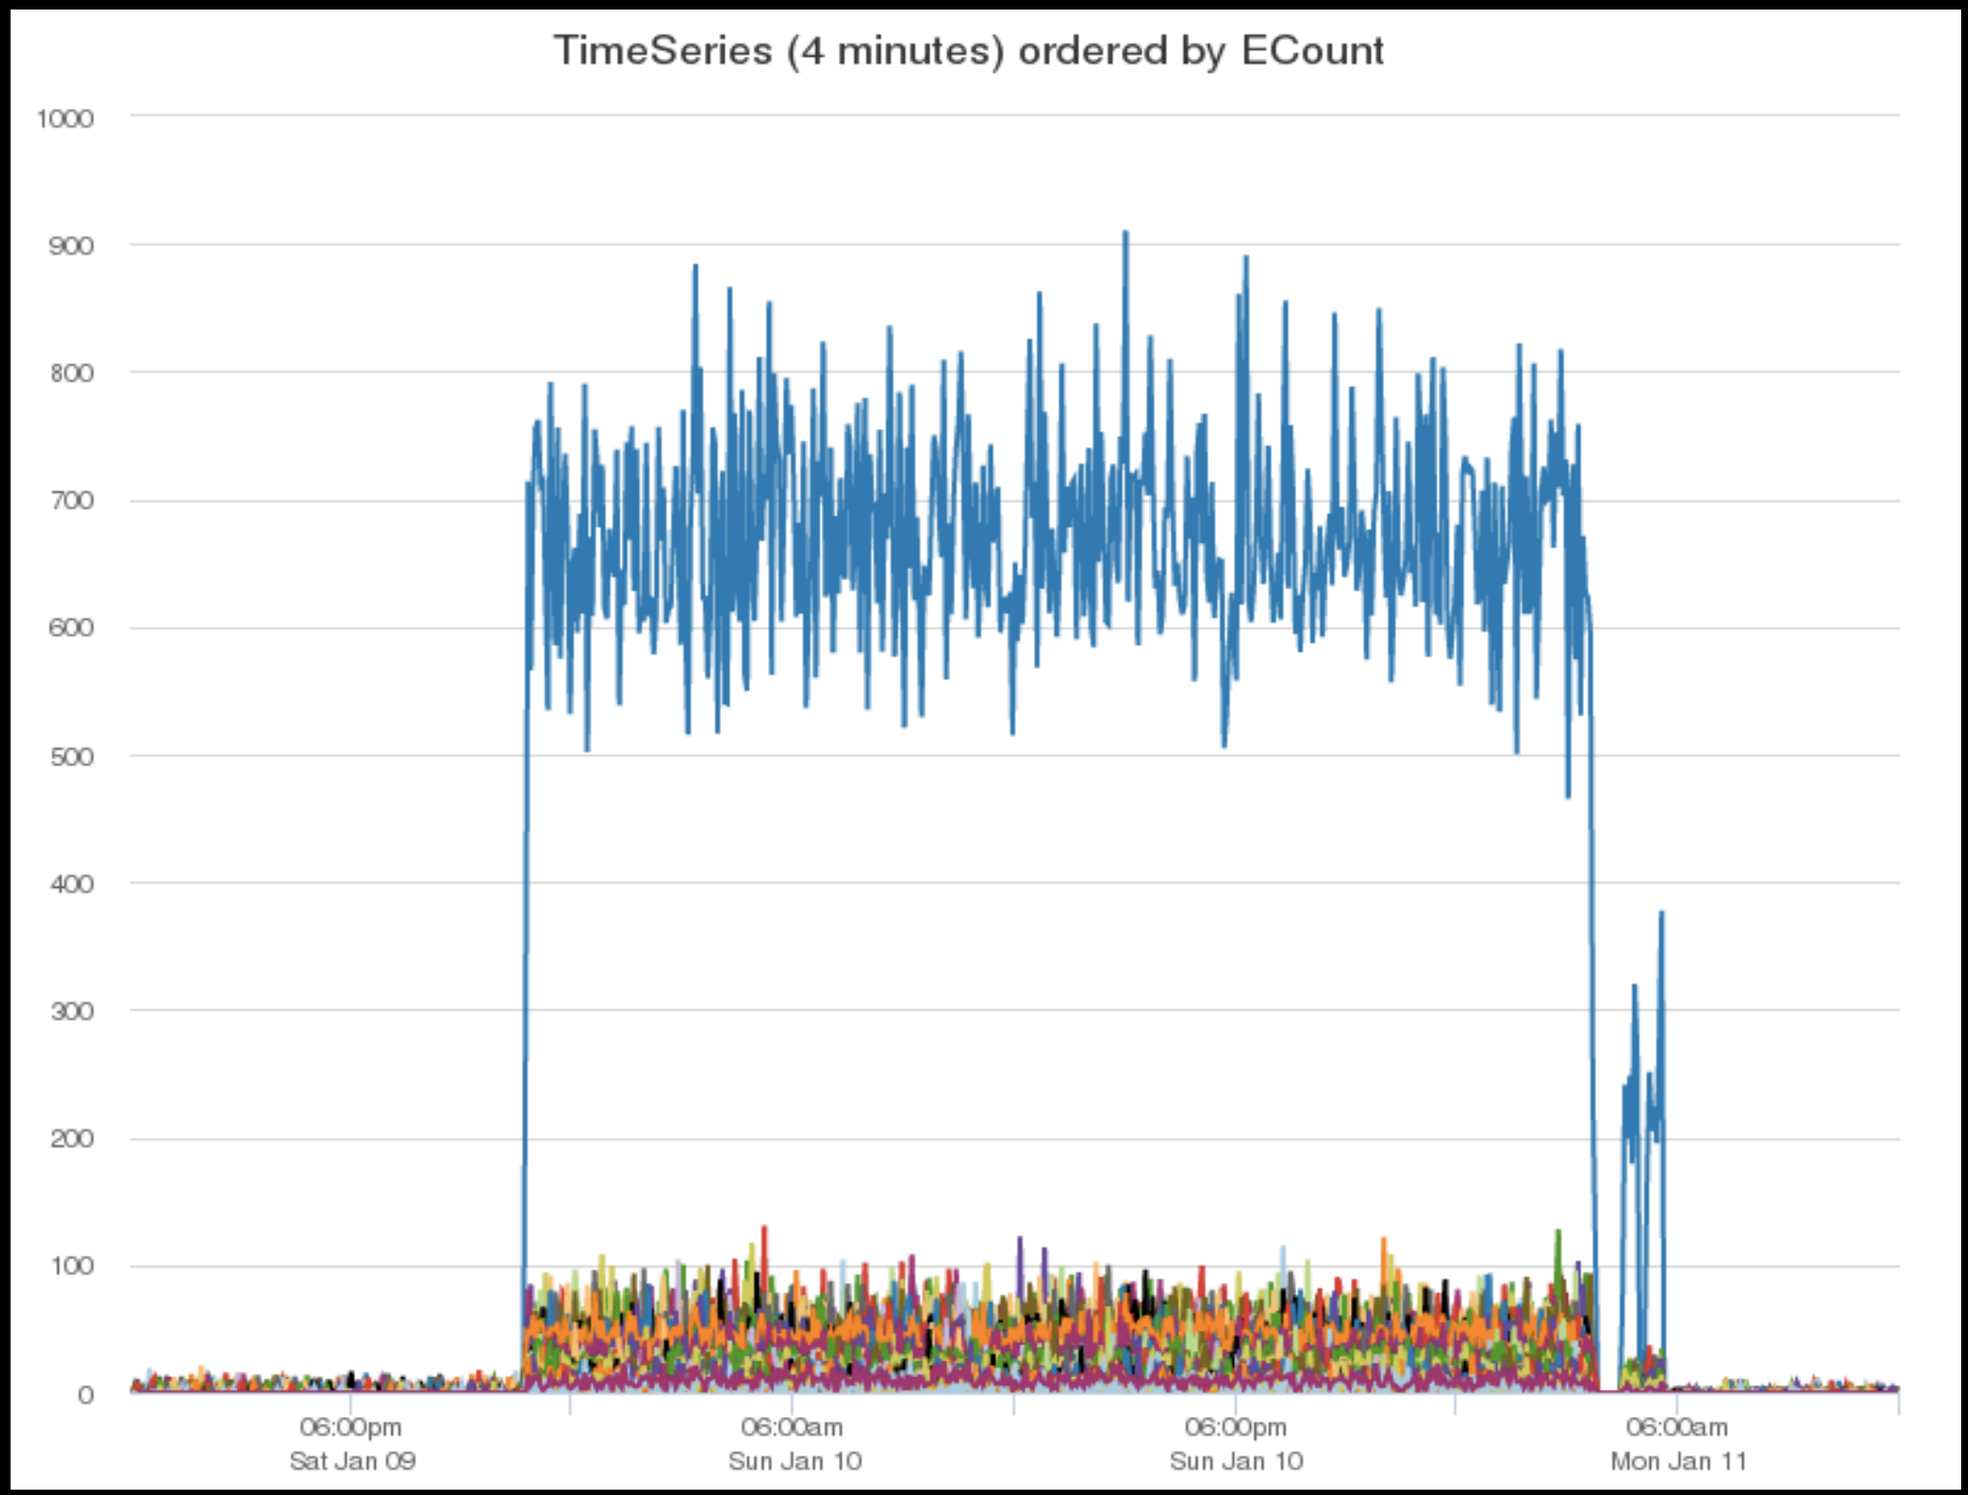

Conclusions

- Fault isolation is actively evolving
- Traceroute + probing approach is quite generic
- Limited by current hardware