“Uptime” at IXPs

- and NIS Directive

Robert Lister

UKNOF 40

27 April 2018 | Manchester

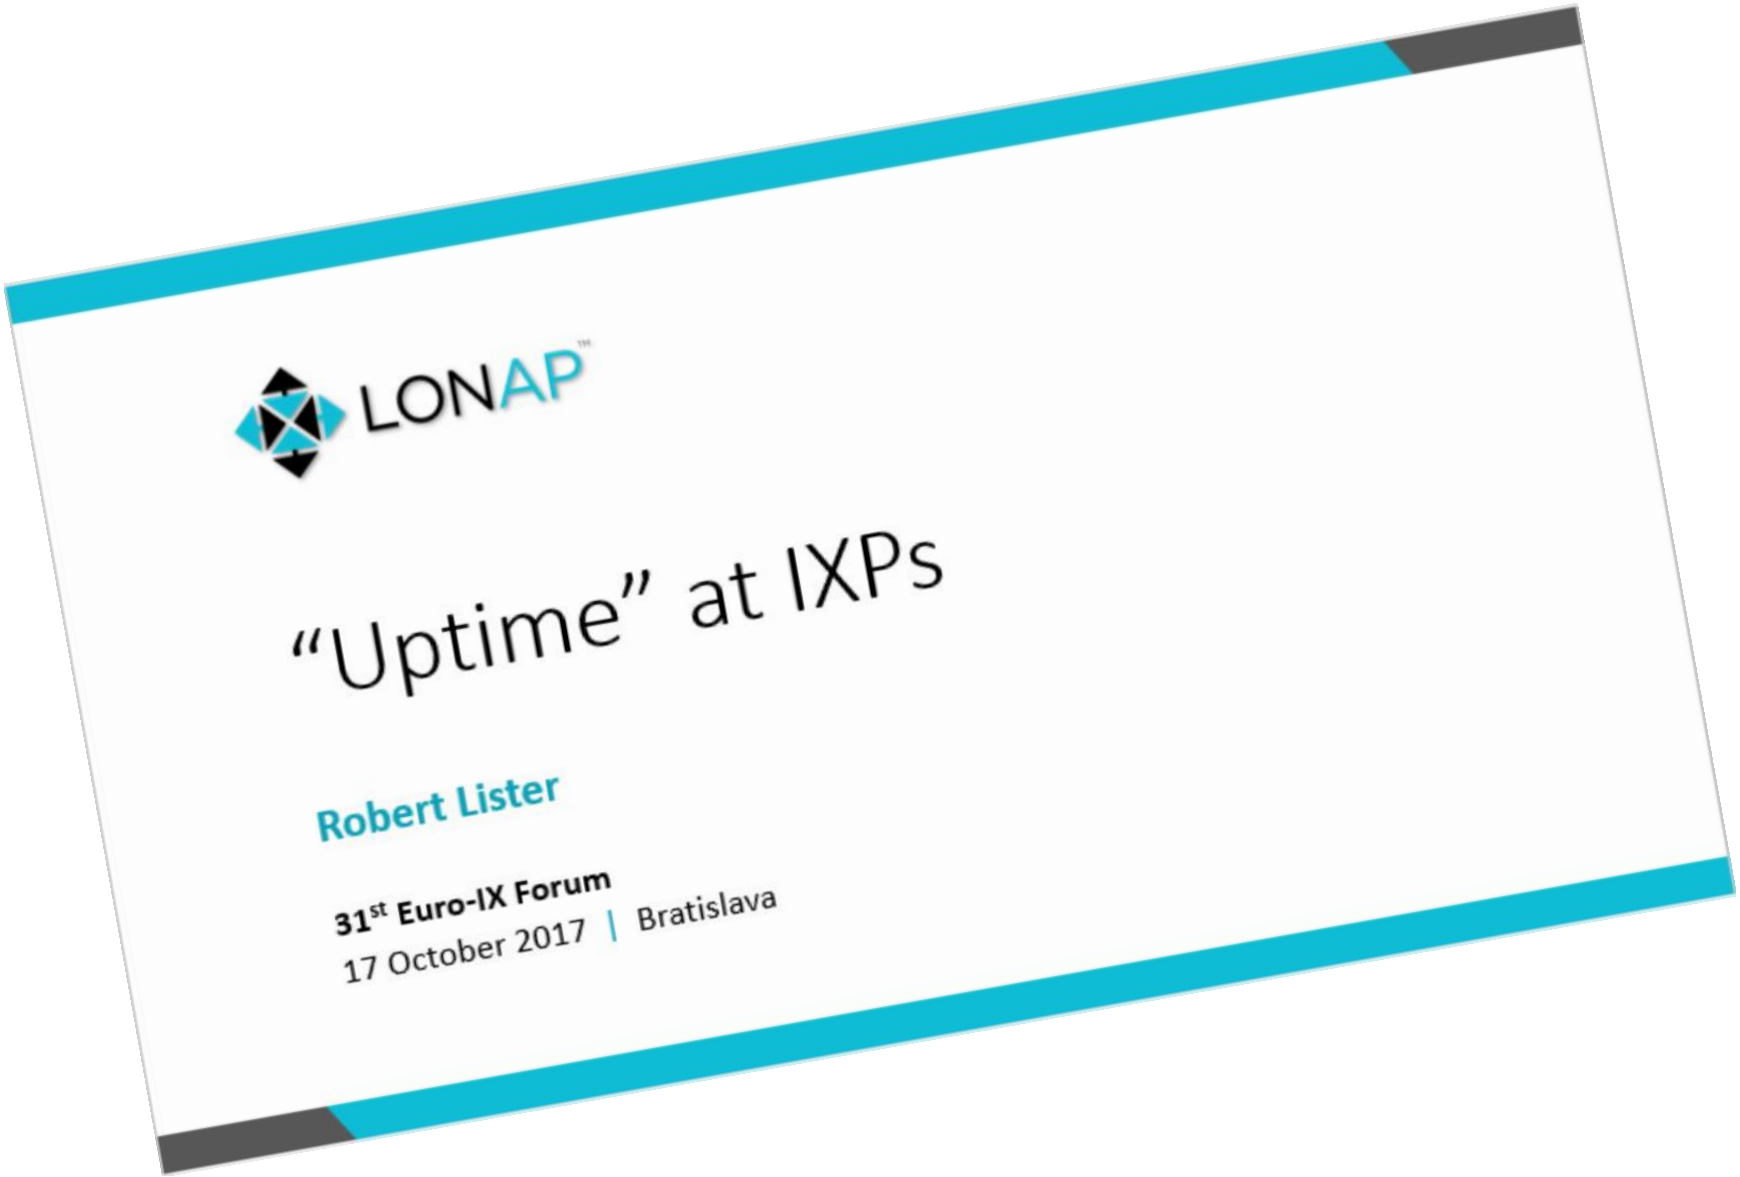
LONAP
"Uptime" at IXPs
Robert Lister
31st Euro-IX Forum
17 October 2017 | Bratislava

LONAP

NIS Directive

- EU Directive on security of Networks and Information Systems

- UK Consultation: (August/Sept 2017): https://www.gov.uk/government/consultations/consultation-on-the-security-of-network-and-information-systems-directive

- https://www.ncsc.gov.uk/guidance/introduction-nis-directive

LONAP™

NIS Directive

- May require **IXPs** to report availability / outage metrics
- For UK, this means OFCOM:

- “Operators who have 50% or more annual market share amongst UK IXP Operators in terms of interconnected autonomous systems,

  Or:

- Who offer interconnectivity to 50% or more of Global Internet routes.”

LONAP™

“High availability”

Availability %	Downtime per year	Downtime per month	Downtime per week	Downtime per day
90% ("one nine")	36.5 days	72 hours	16.8 hours	2.4 hours
95% ("one and a half nines")	18.25 days	36 hours	8.4 hours	1.2 hours
97%	10.96 days	21.6 hours	5.04 hours	43.2 minutes
98%	7.30 days	14.4 hours	3.36 hours	28.8 minutes
99% ("two nines")	3.65 days	7.20 hours	1.68 hours	14.4 minutes
99.5% ("two and a half nines")	1.83 days	3.60 hours	50.4 minutes	7.2 minutes
99.8%	17.52 hours	86.23 minutes	20.16 minutes	2.88 minutes
99.9% ("three nines")	8.76 hours	43.8 minutes	10.1 minutes	1.44 minutes
99.95% ("three and a half nines")	4.38 hours	21.56 minutes	5.04 minutes	43.2 seconds
99.99% ("four nines")	52.56 minutes	4.38 minutes	1.01 minutes	8.64 seconds
99.995% ("four and a half nines")	26.28 minutes	2.16 minutes	30.24 seconds	4.32 seconds
99.999% ("five nines")	5.26 minutes	25.9 seconds	6.05 seconds	864.3 milliseconds
99.9999% ("six nines")	31.5 seconds	2.59 seconds	604.8 milliseconds	86.4 milliseconds
99.99999% ("seven nines")	3.15 seconds	262.97 milliseconds	60.48 milliseconds	8.64 milliseconds
99.999999% ("eight nines")	315.569 milliseconds	26.297 milliseconds	6.048 milliseconds	0.864 milliseconds
99.9999999% ("nine nines")	31.5569 milliseconds	2.6297 milliseconds	0.6048 milliseconds	0.0864 milliseconds

LONAP™

Source: https://en.wikipedia.org/wiki/High_availability

99.99(9)% uptime?

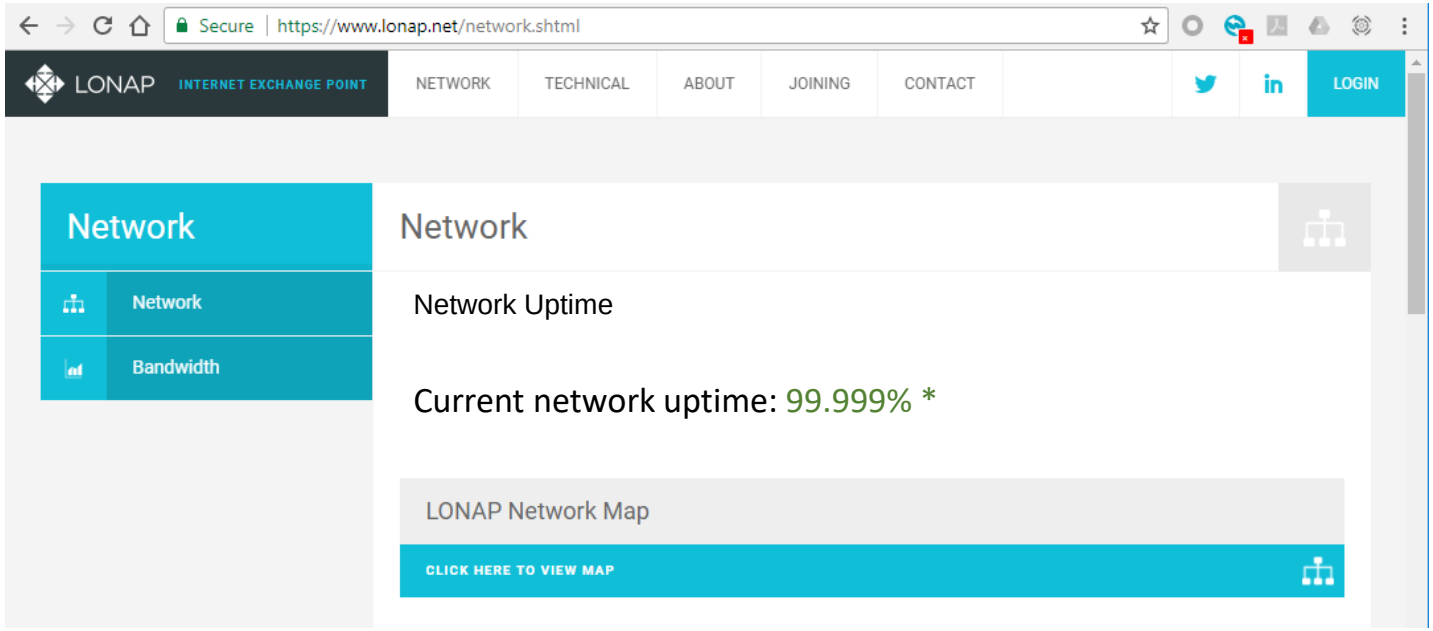

LONAP™

99.99(9)% uptime?

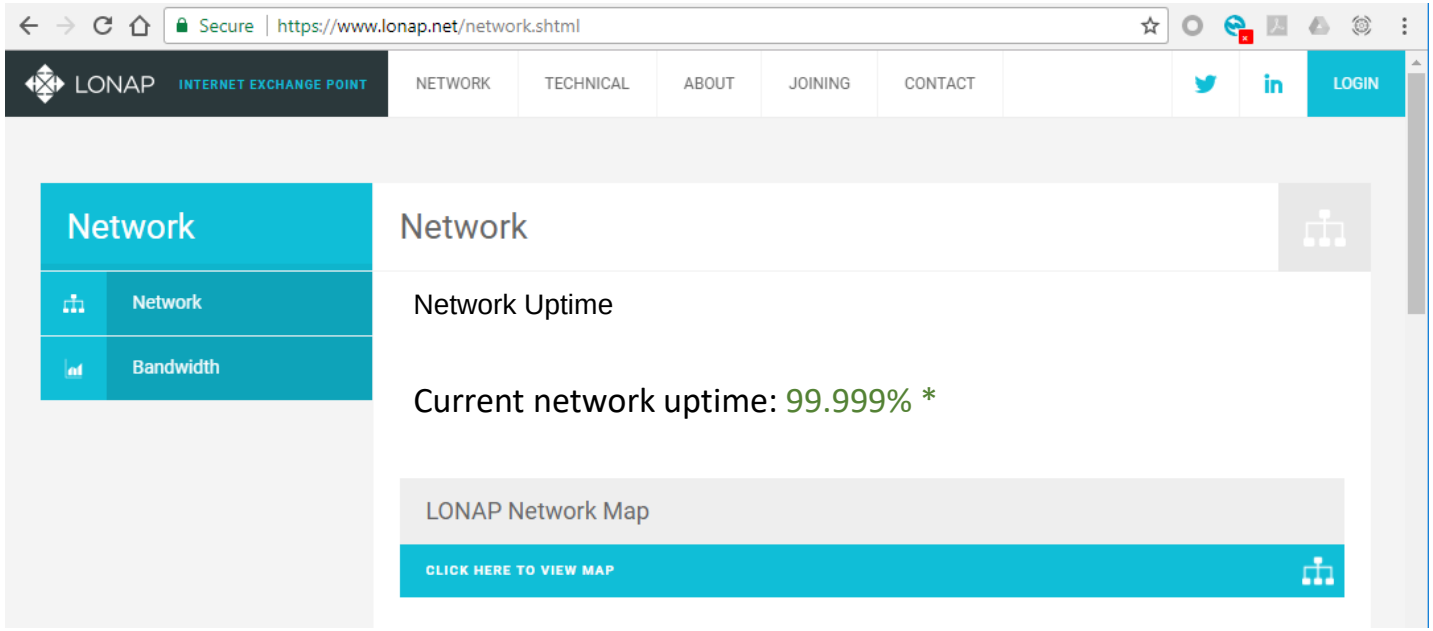

- *9 out of 10 cats local pref our prefixes. The value of your pings may go down as well as up.*
- *We reserve the right to replace lost packets with equivalent size packets at our discretion.*
- *Not to scale. Not actual web site.*
- *Due to rounding, numbers presented may not add up precisely to the totals provided and percentages may not precisely reflect the absolute figures. Figures were correct at time we made them up.*
- *Subject to National Rail Conditions of Travel. Packets valid via any reasonable route.*
- *Contents may settle during shipping.*

LONAP™

Determine “up” at an IXP

member	ping?
5.57.80.1	✓
5.57.80.2	✓
5.57.80.3	✓
5.57.80.4	✓
5.57.80.5	✓
…etc…	
5.57.80.xx	✓

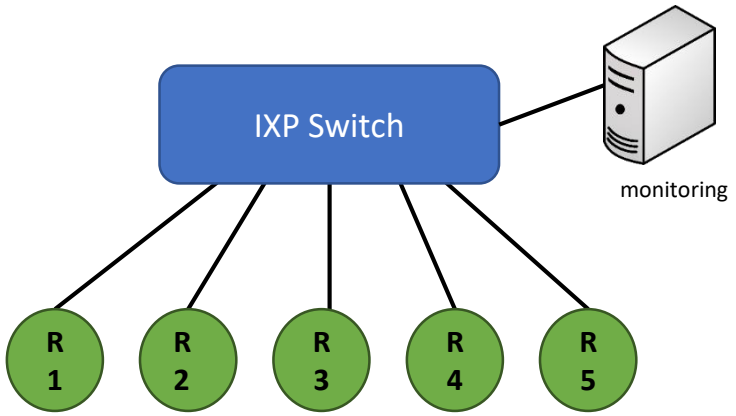

= “100% up”

LONAP™

Ping all the things…

member	ping	ping	ping	ping	ping	ping	ping	ping	ping	ping	ping	Available %
5.57.80.1	✓	✓	✓	✓	✓	✓	✓	✓	✓	… lots more columns …	✓	100%
5.57.80.2	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	100%
5.57.80.3	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	100%
5.57.80.4	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	100%
5.57.80.5	✓	✓	✗	✗	✗	✗	✗	✓	✓		✓	99.65%

Example:

- In 24 hours = 1440 minutes.
- -5 minutes downtime = 1435 (99.652%)
- It would more likely be calculated in seconds: (86400 – 300 = 99.652%)

LONAP™

Pinging members can suck…

member	ping	ping	ping	ping	ping	ping	ping	ping	ping	ping
5.57.80.1	✓	✓	✗	✗	✗	✓	✗	✓	✗	✓
5.57.80.2	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
5.57.80.3	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
5.57.80.4	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓
5.57.80.5	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗

- Some members may have busy routers (high latency/packet loss)
- Some do not reply to ping
- Might miss shorter outages between pings
- Latency is an interesting stat to monitor

LONAP™

It can get ……. *messy*

member	ping	ping	ping	ping	ping	ping	ping	ping	ping	ping
5.57.80.1	✓	✓	✗	✗	✗	✓	✗	✓	✗	✓
5.57.80.2	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
5.57.80.3	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
5.57.80.4	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓
5.57.80.5	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗

- IXP Manager option:

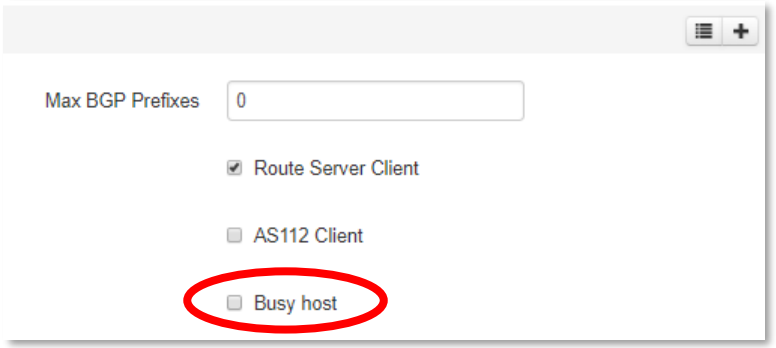

LONAP™

Correlate pings with other pings!

member	ping	ping	ping	ping	ping	ping	ping	ping	ping	ping
5.57.80.12	✓	✓	✓	✓	✗	✗	✗	✗	✗	✓
5.57.80.52	✓	✓	✓	✓	✗	✗	✗	✗	✗	✓
5.57.80.48	✓	✓	✓	✓	✗	✗	✗	✗	✗	✓
5.57.80.76	✓	✓	✓	✓	✗	✗	✗	✗	✗	✓
5.57.80.91	✓	✓	✓	✓	✗	✗	✗	✗	✗	✓

- Pinging a single host is limited *by itself*: **more useful if we correlate**
- Multiple members unreachable in the same interval.
  - May indicate an outage?

LONAP™

Correlate other monitoring data

member	ping	BGP	RS1	RS2	Port	ARP	traffic	errors	…
5.57.80.12	✘	✓	✓	✓	✓	✓	50%	0	
5.57.80.52	✘	✘	✘	✘	✘	✘	0%	0	
5.57.80.48	✘ 7/10	✓	✓	✓	✓	✓	99%	5068	
5.57.80.76	✓	✓	✓	✓	✓	✓	38%	0	
5.57.80.91	✓	✘	✘	✘	✓	✓	0%	0	

- Correlating with *other monitoring* gives us more insight
- This is useful for monitoring ☺
- Makes a “single metric” calculation **complex** ☹
- It is both up *and* down? Wait a bit…

```
# My clever alert
correlation script 1.0

if ($port_down) {
  if (…) {

…lots of twisty code

 }
}

$uptime = do_magic()

# 2002-08-10: should
# probably rewrite this
# bit sometime…

# 2018-01-28: LOL!

@PORTS = get_snmp_voodoo()
```

LONAP

Path availability

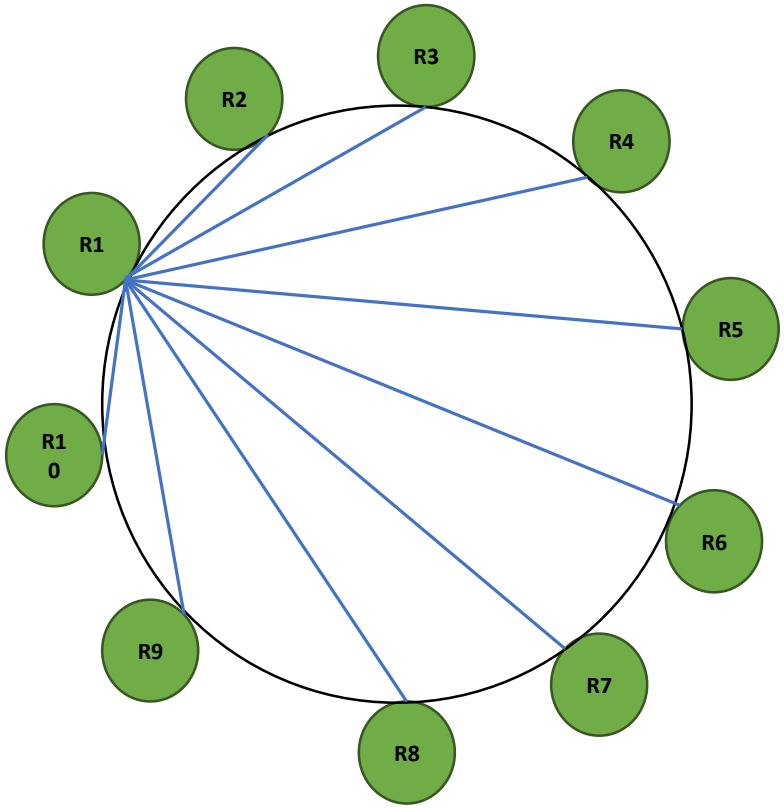

LONAP™

Path availability

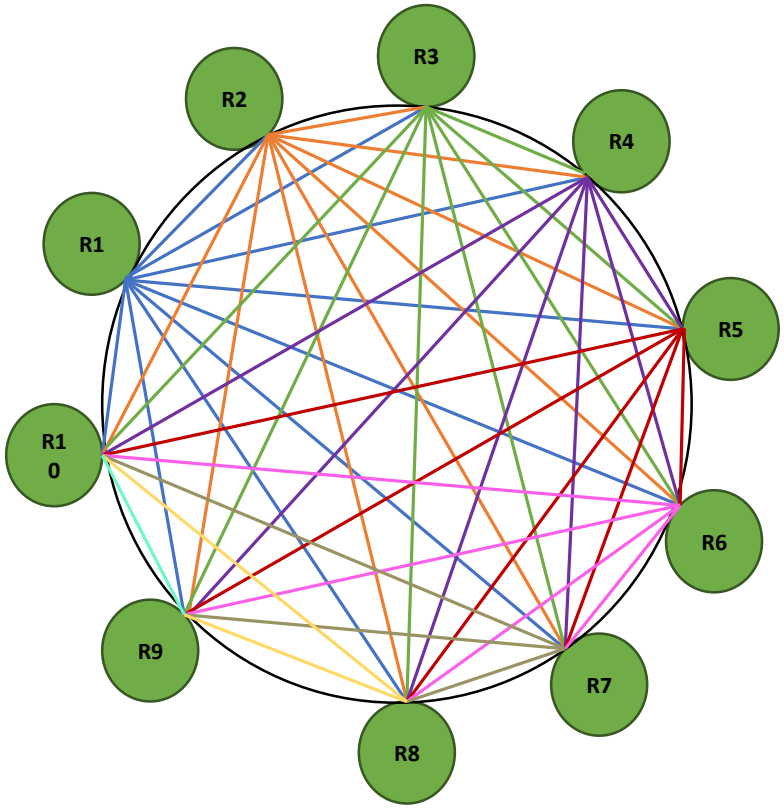

possible paths =

`n * (n-1) / 2`

`10 * (10-1) / 2 =` **45**

(45 paths available = 100%)

We consider every path, whether or not peering exists

ASNs don't peer with themselves.

yes, this slide took forever to draw…

LONAP™

Exchange topology

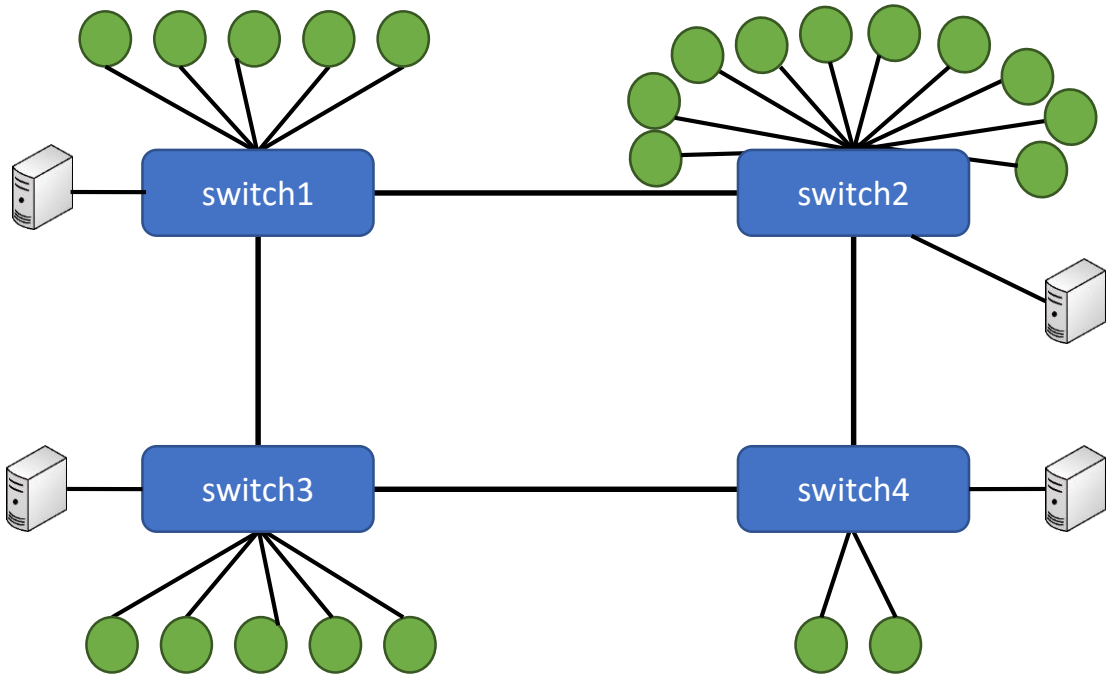

LONAP™

Exchange topology

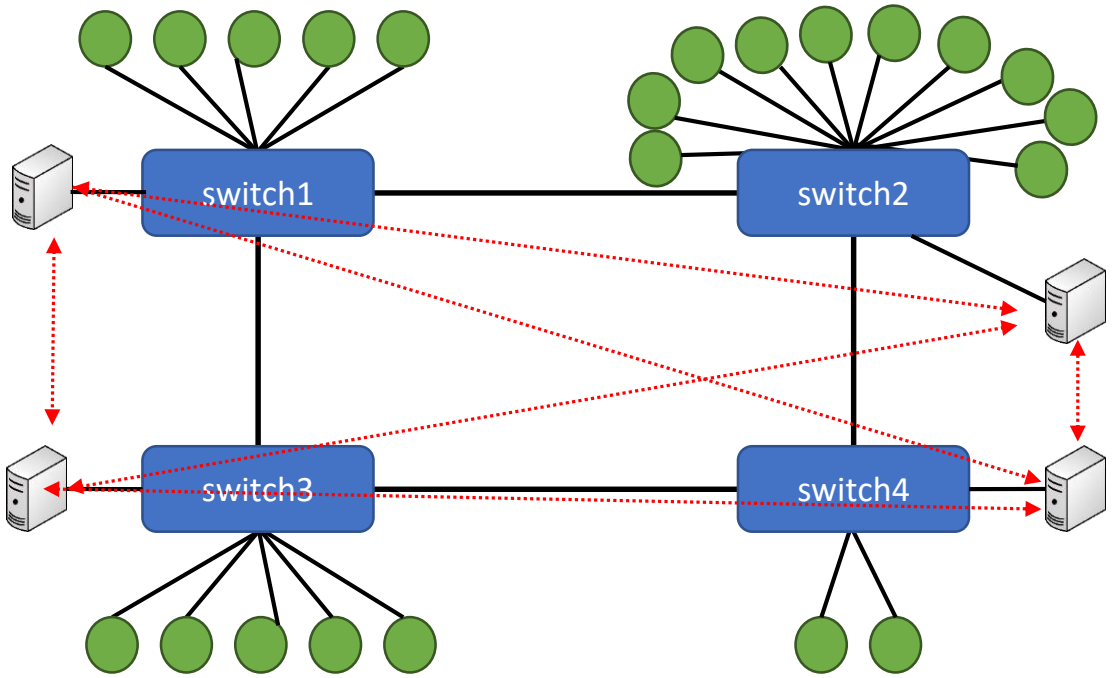

LONAP

Calculating path availability

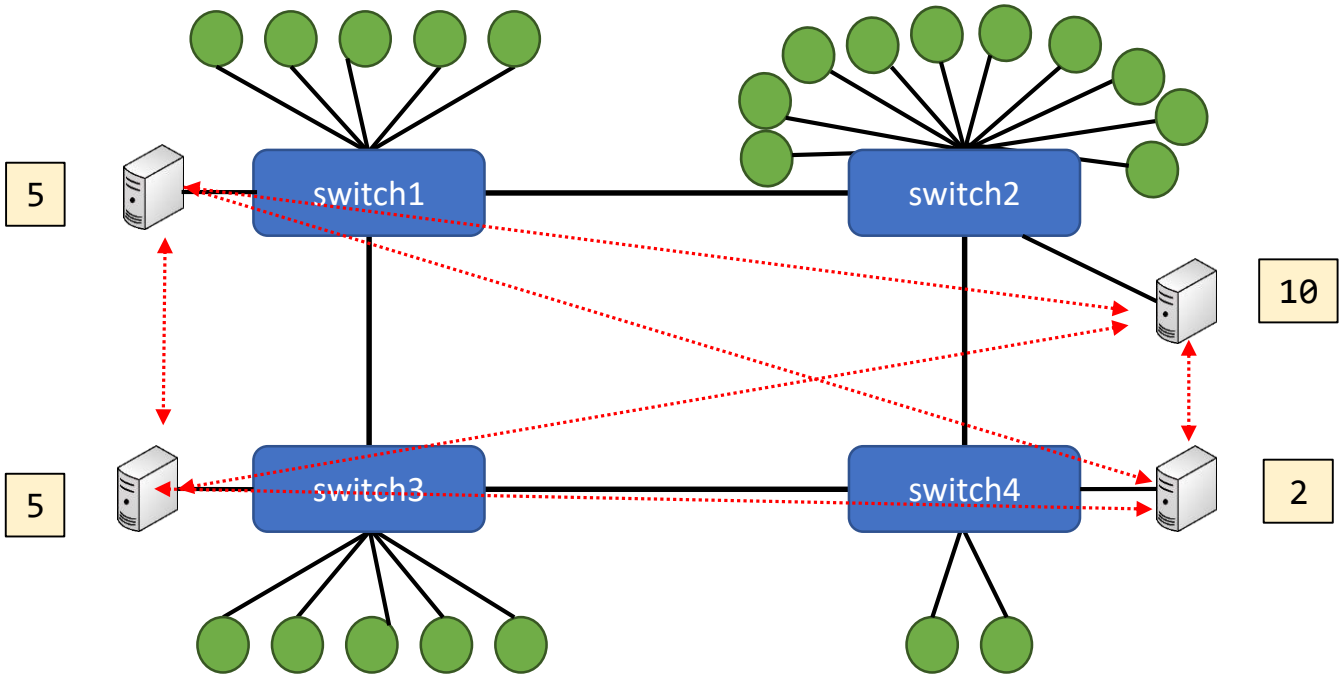

LONAP

Calculating path availability

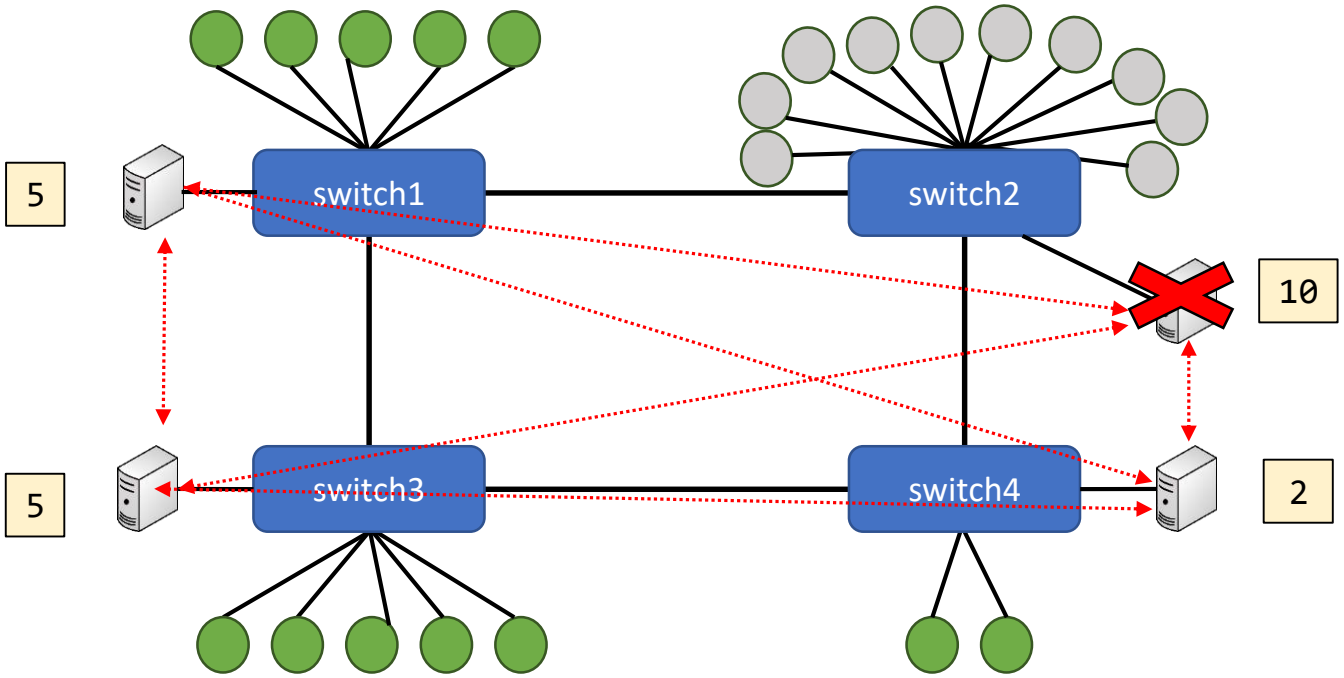

LONAP

Calculating path availability

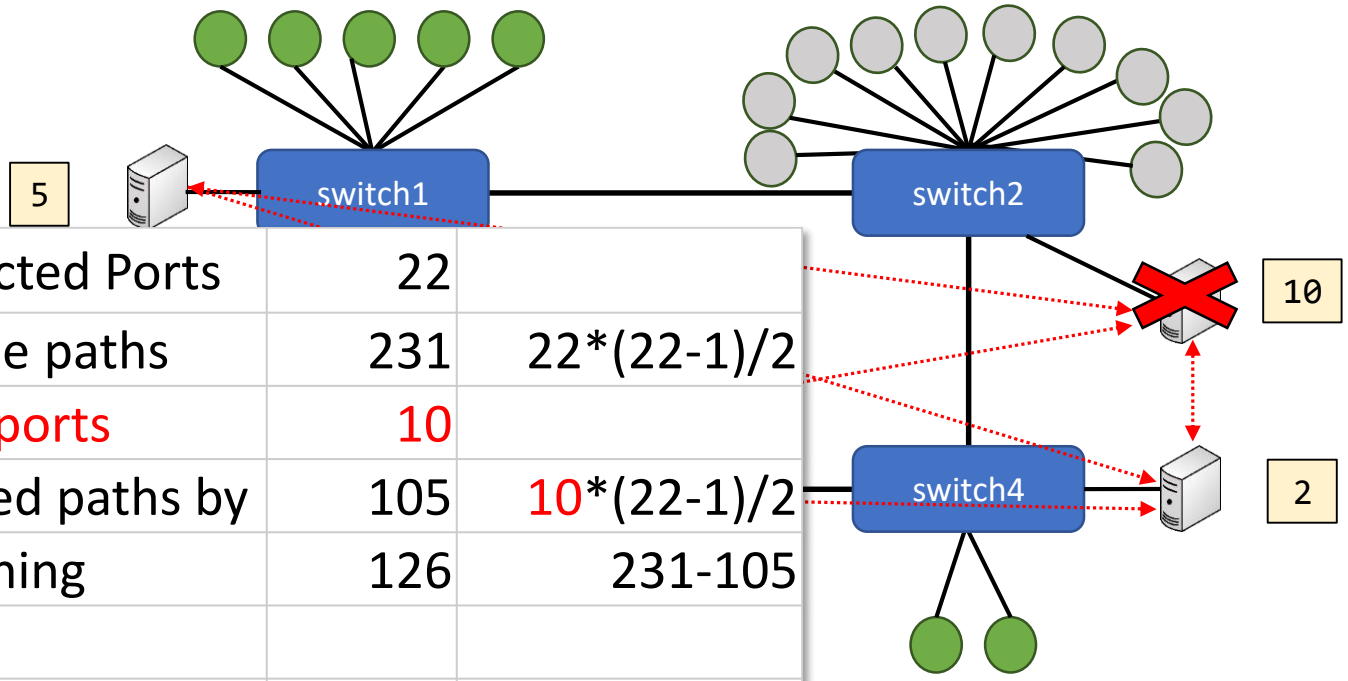

Connected Ports	22	
Possible paths	231	22*(22-1)/2
Down ports	10	
Reduced paths by	105	10*(22-1)/2
Remaining	126	231-105
Path Availability	54.55%	

LONAP

Calculating path availability

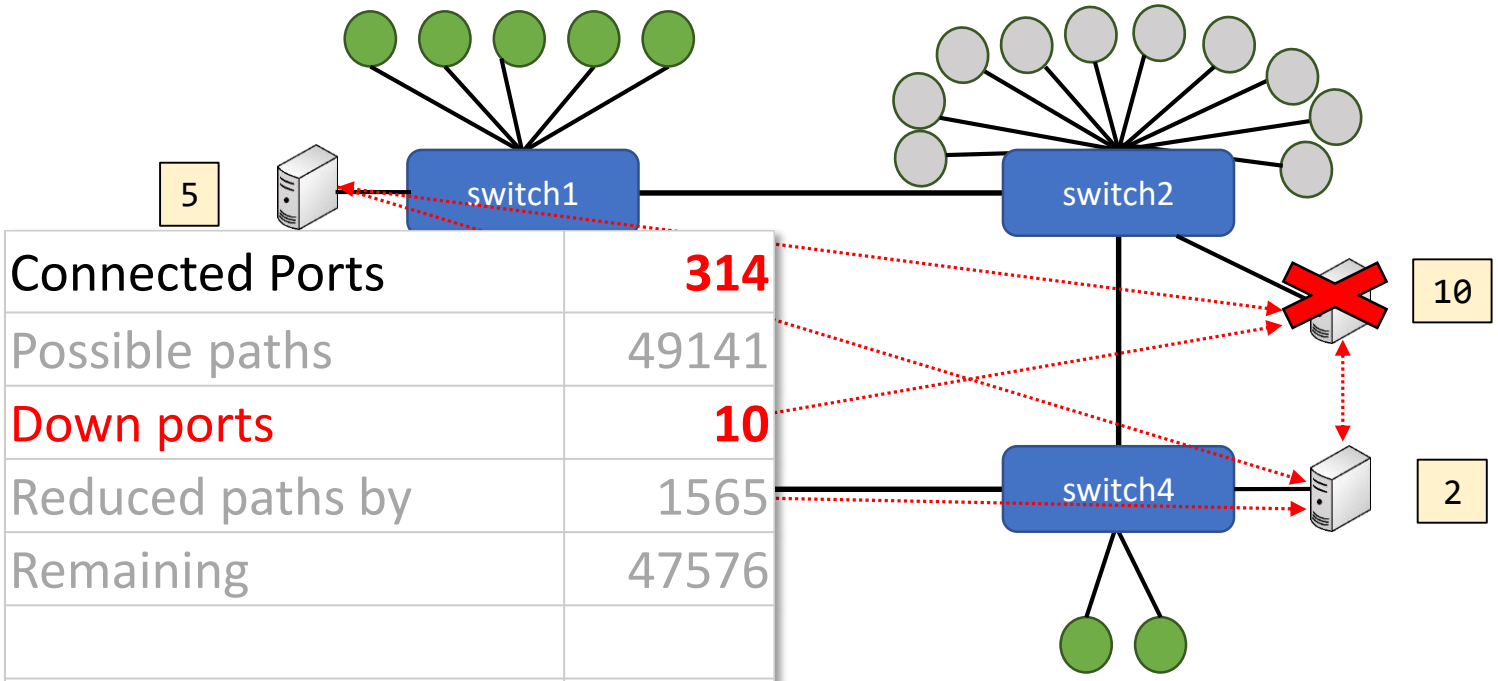

Connected Ports	**314**
Possible paths	49141
Down ports	**10**
Reduced paths by	1565
Remaining	47576
Path Availability	96.82%

LONAP

..another way to do it – by **port capacity?**

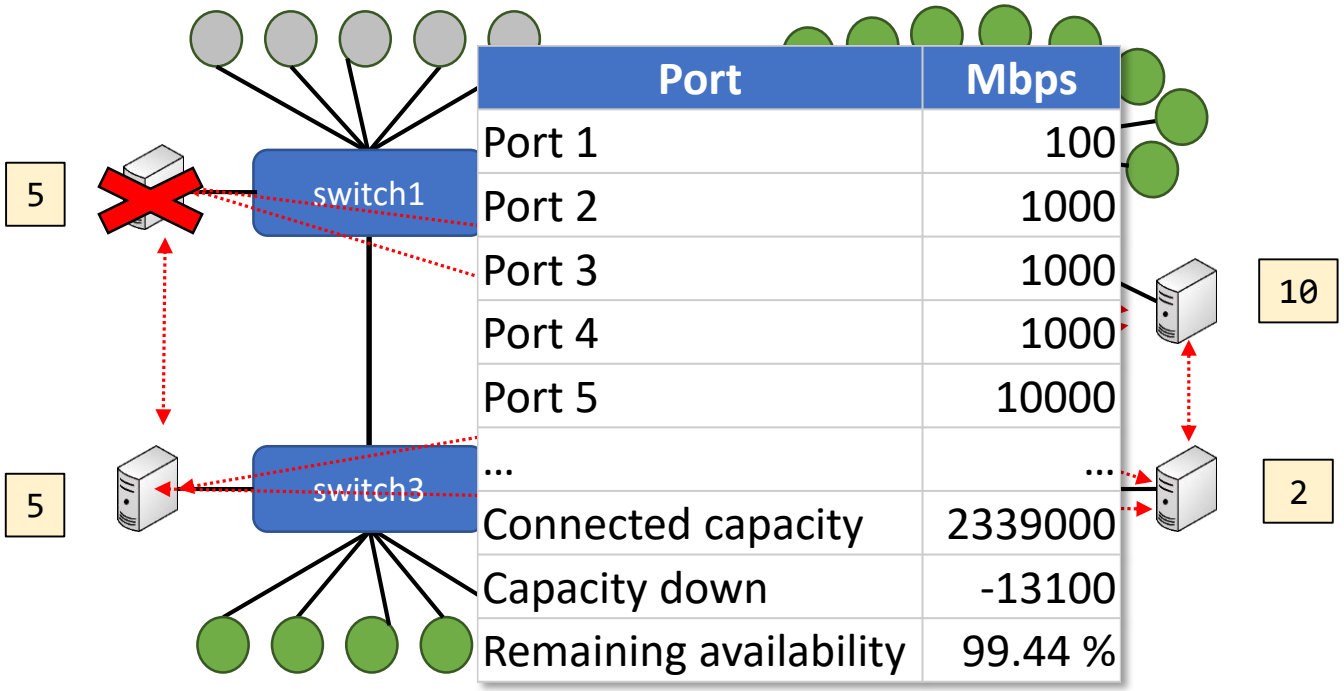

Port	Mbps
Port 1	100
Port 2	1000
Port 3	1000
Port 4	1000
Port 5	10000
…	…
Connected capacity	2339000
Capacity down	-13100
Remaining availability	99.44 %

LONAP

..another way to do it – by **port capacity?**

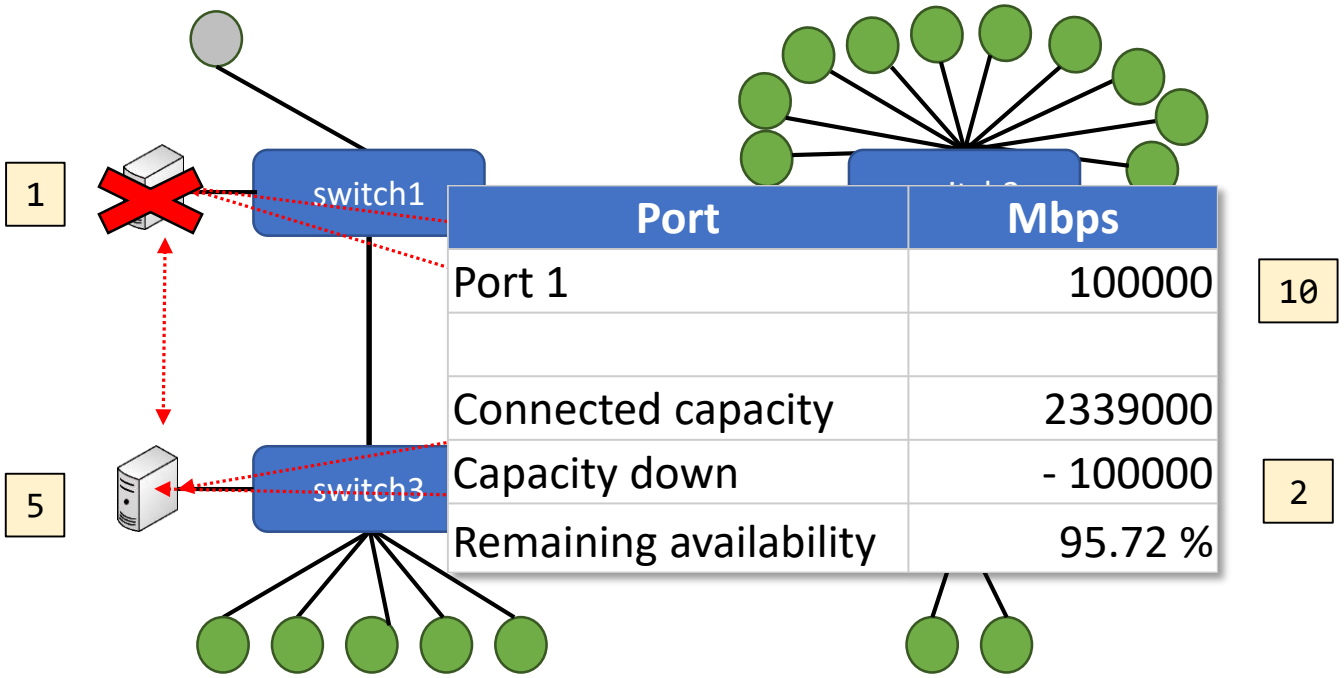

LONAP

…or use the switches themselves?

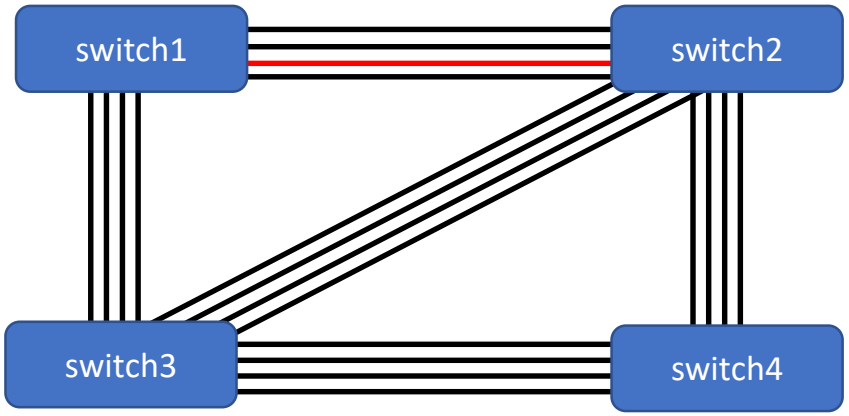

- No longer just a flat layer 2 network. The devices are layer 3.
  - Every core link is an IP, point-to-point link
  - …we could monitor these to work out “core availability”
  - Maybe take into account **traffic impact** (down link may have no noticeable impact)

LONAP™

Is it a **useful** metric?

- Do we exclude things like **maintenance**?
- Exclude other factors "**outside our control?**"
- Is that realistic?

- Try not to obsess about the number!

100% 99.99% 100% 100% 100% 100% 100% 100% 100% 100% 100% 100%% 100% 100% 100% 100% 100% 100% 100% 100%

LONAP™

What LONAP members said…

- “Your job is to move packets. Just monitor ingress and egress packets”
- “Don’t spend a lot of effort creating this metric.”
- “Just focus on running a reliable service. Don’t break it.”
- Use SFLOW to detect problems (find increased TCP SYN)”
- “Use whatever metric internally if it helps. Probably not useful to publish it.”
- “You need more pictures of cats.”

LONAP™

What other EURO-IX IXPs said…

- “We tried and gave up.”
- “It’s too complicated to create a reliable number”
- “We do a complex calculation to create availability metrics”
- Should we try to develop some standard metrics?

LONAP

Thoughts?