CSE 550: Systems for all

Au 2022

Ratul Mahajan

Routing

Find path between any source-destination pair

- Needed because not all pairs are connected directly

Intra-domain routing: Find short paths (OSPF, ISIS)

Inter-domain routing: Find policy-compliant paths (BGP)

Impact of policy in BGP

Performance

Convergence

Security

BGP performance

Each party selects routes to suit its own interests

- e.g, shortest path in ISP

What path will be chosen for A2→B1 and B1→A2?

- What is the best path?

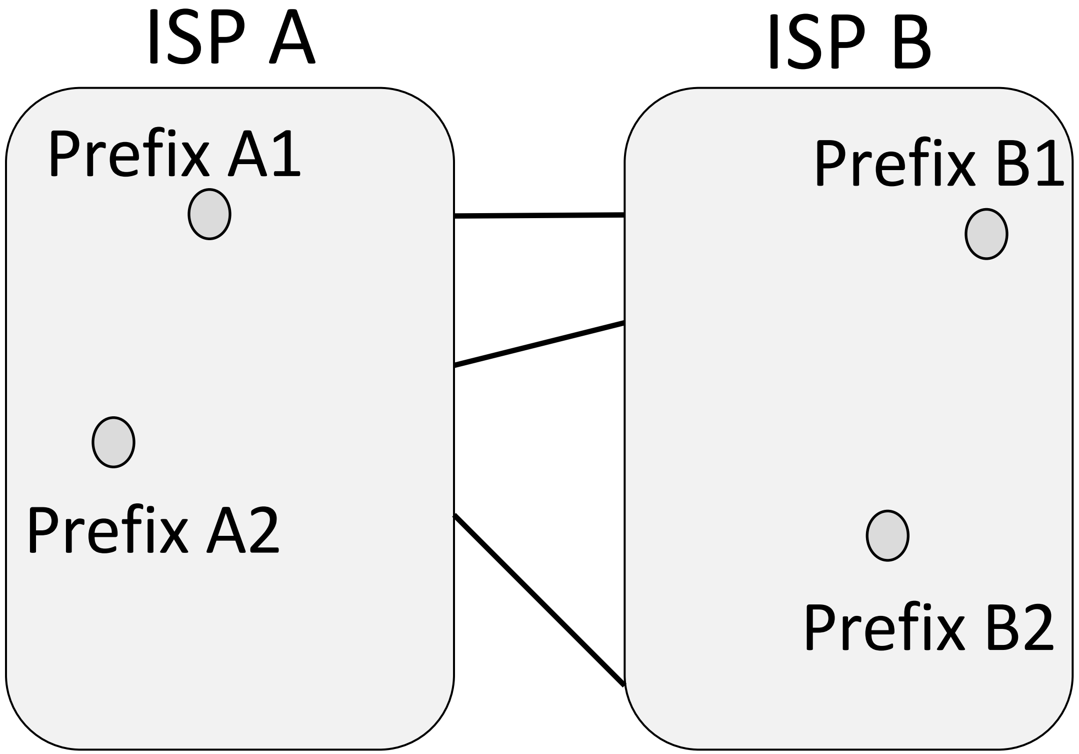

BGP performance

Selected paths are longer than overall shortest path

- And asymmetric too!

Consequence of independent goals and decisions (not hierarchy and info hiding)

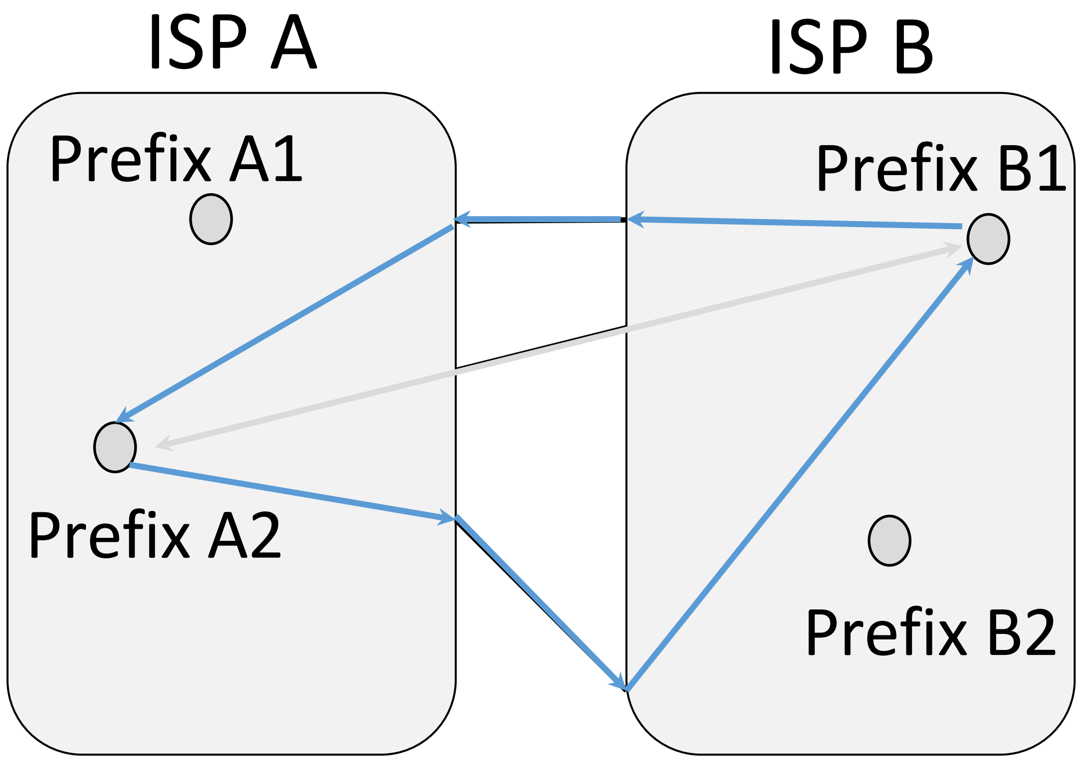

BGP slow convergence

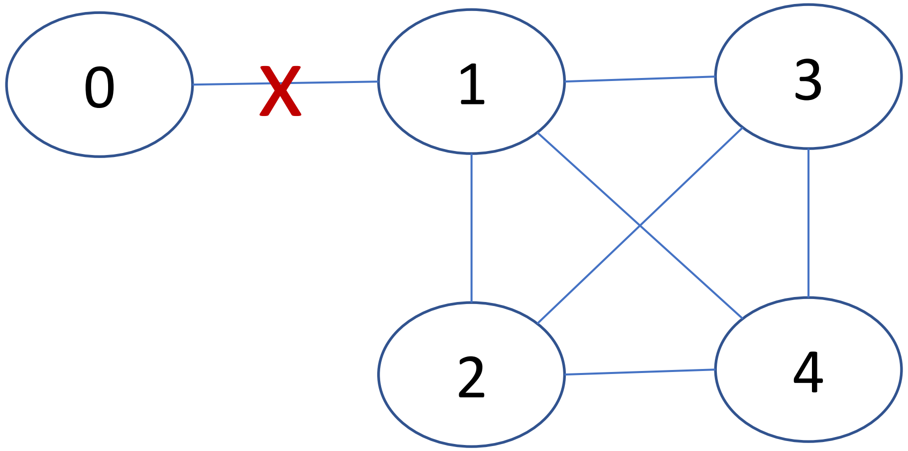

[1, 0]

[4, 1, 0]
[2, 1, 0]

[1, 0]

[3, 1, 0]
[4, 1, 0]

[1, 0]

[2, 1, 0]
[3, 1, 0]

BGP slow convergence

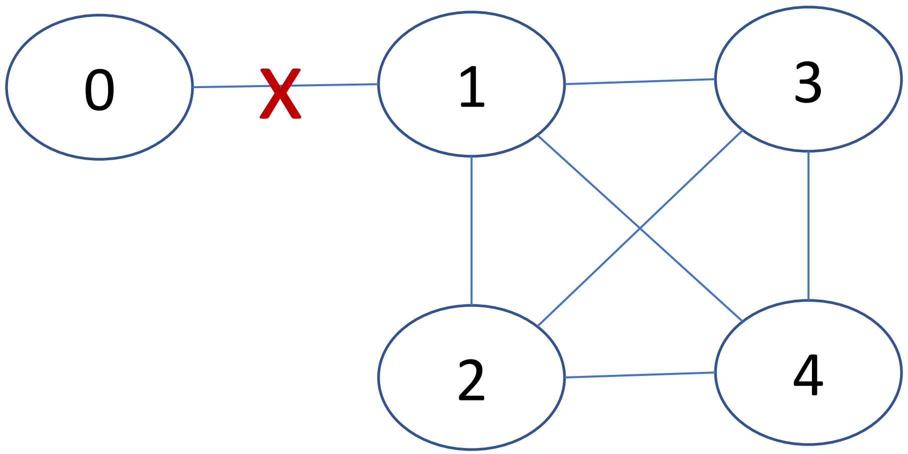

[4, 1, 0]

[2, 1, 0]

[3, 1, 0]

[4, 1, 0]

[2, 1, 0]

[3, 1, 0]

BGP slow convergence

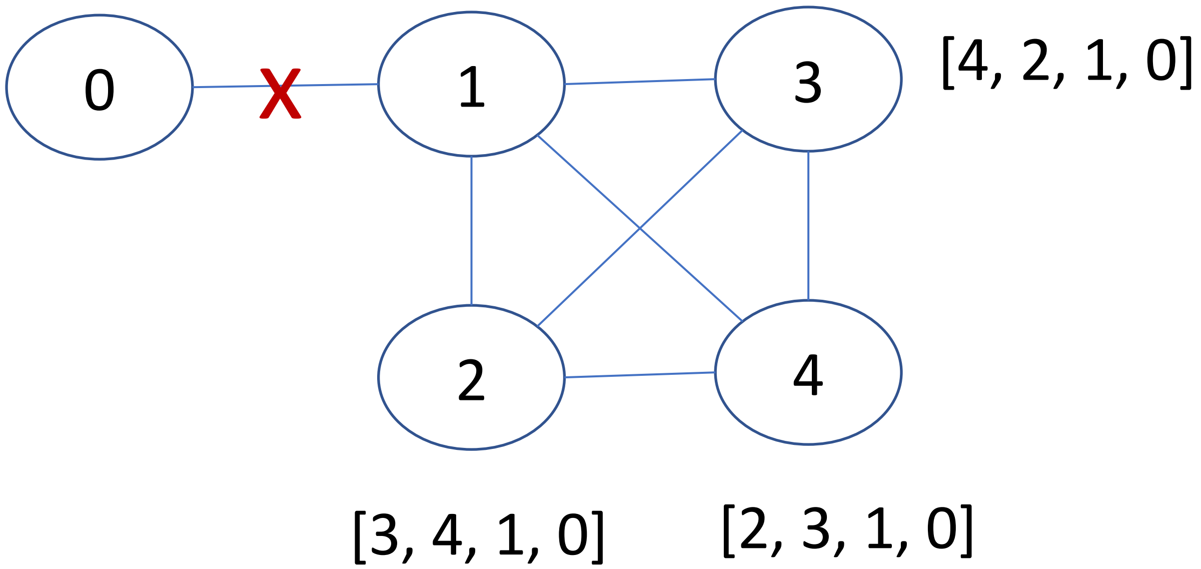

BGP “bad gadget”: Non-convergence

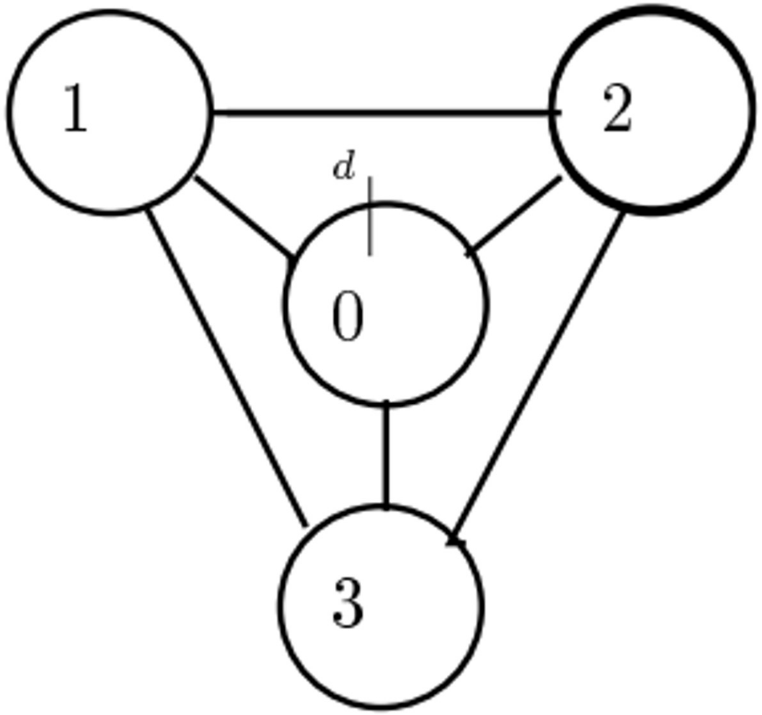

[2, 0] > [0] > [2, 3, 0]

[3, 0] > [0] > [3, 1, 0]

[1, 0] > [0] > [1, 2, 0]

BGP insecurity

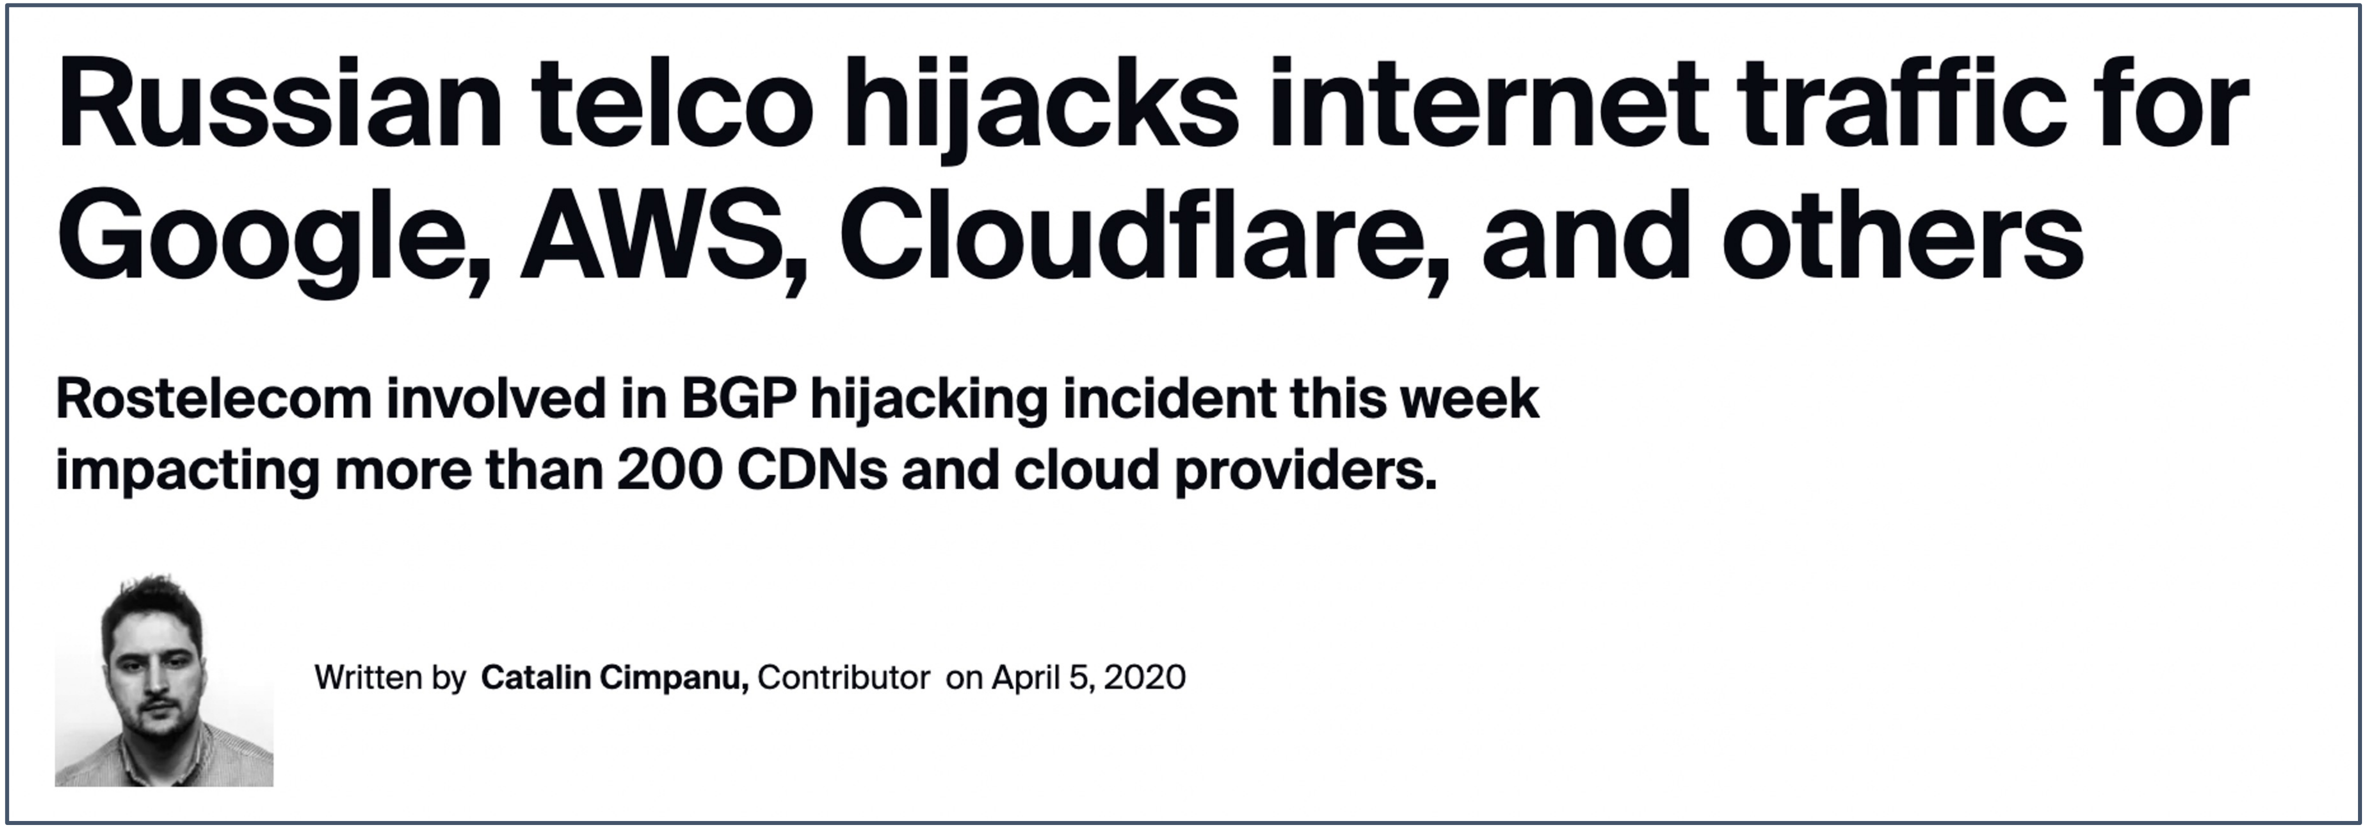

Russian telco hijacks internet traffic for Google, AWS, Cloudflare, and others

Rostelecom involved in BGP hijacking incident this week impacting more than 200 CDNs and cloud providers.

Written by **Catalin Cimpanu,** Contributor on April 5, 2020

BGP insecurity

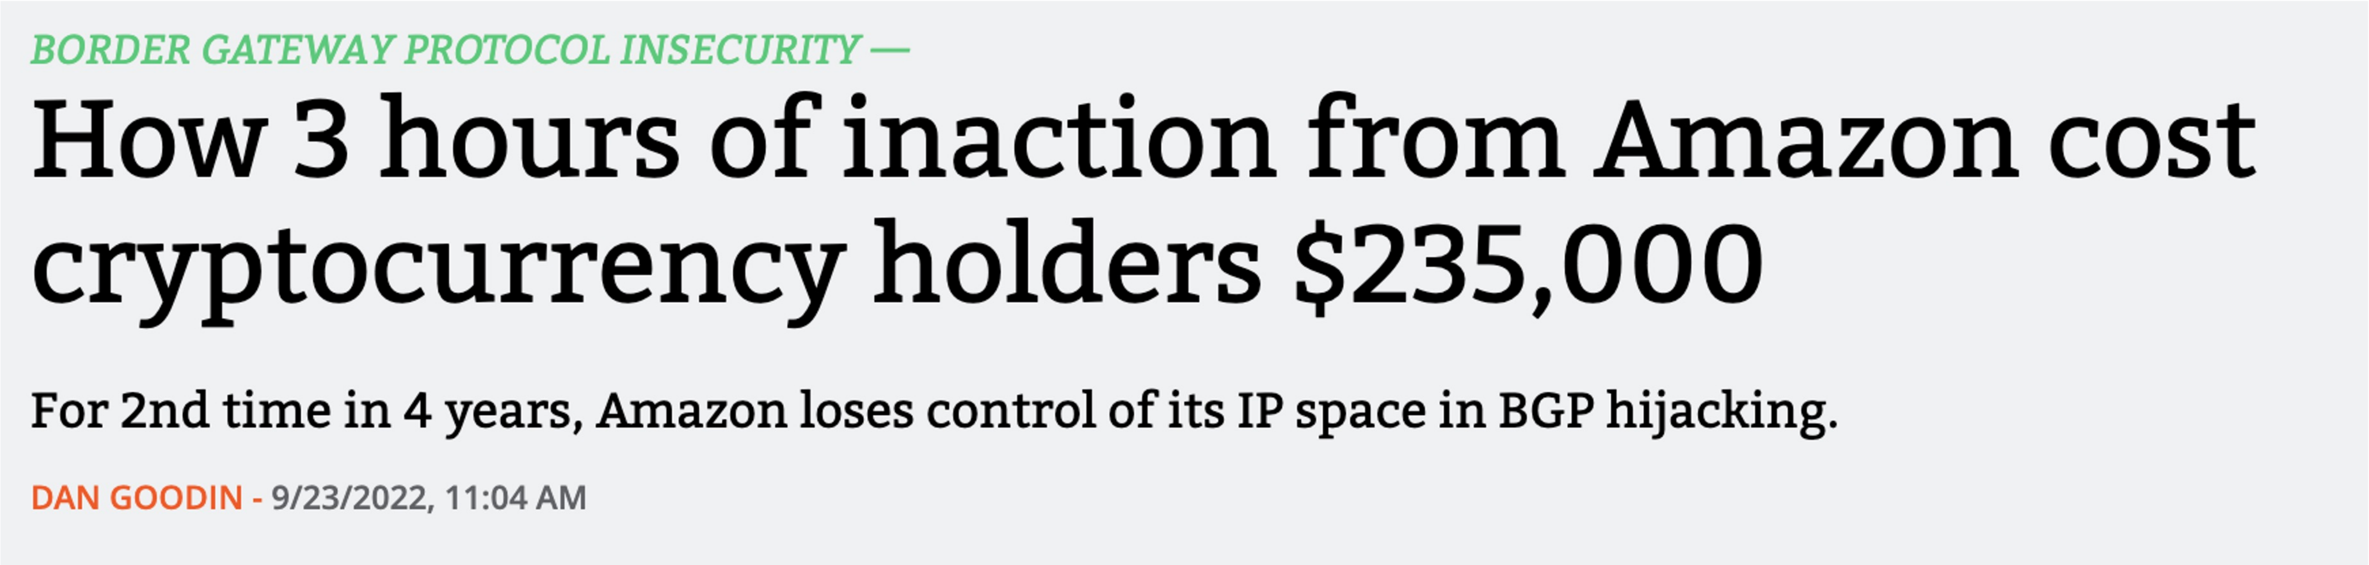

BORDER GATEWAY PROTOCOL INSECURITY —

How 3 hours of inaction from Amazon cost cryptocurrency holders $235,000

For 2nd time in 4 years, Amazon loses control of its IP space in BGP hijacking.

DAN GOODIN - 9/23/2022, 11:04 AM

BGP security: Why is it hard?

No authoritative database of who owns what

- Though getting closer with RPKI

No authoritative database of who is connected to whom

- Peering DBs exist but cannot be relied upon for being up-to-date

No authoritative database of policy

- May not ever happen because of information sensitivity

FB Oct 2021 outage: Murder, suicide, and obstruction

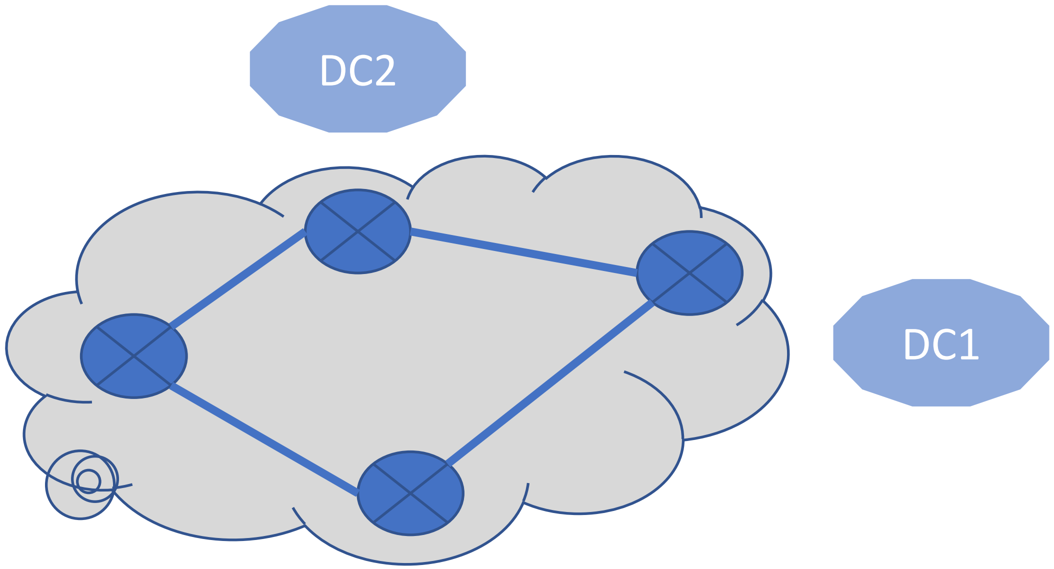

FB Oct 2021 outage: Murder, suicide, and obstruction

Normal operation	Murder	Suicide	Obstruction
FB has a global backbone that connects its DCs and a distributed DNS infra. DNS servers measure “distance” to different DCs. DNS offers “close” DC prefixes to users and withdraws (from BGP) for unreachable DCs	An engineer or a script sends a bad command to backbone routers. Audit tool fails to detect the bad command. All DCs are disconnected from the backbone.	DNS can no longer any DC (since they are all disconnected). DNS withdraws many prefixes from BGP. The prefixes cover the DNS infra as well, so DNS makes itself unreachable as well.	Service restoration requires manual intervention (since nothing is reachable). Physical access requires authorization that turn depends on the same DNS. Takes multiple hours to override systems and gain access to the equipment.

Lessons?

Over to Dixon and Winston