

# Software-Defined Networking

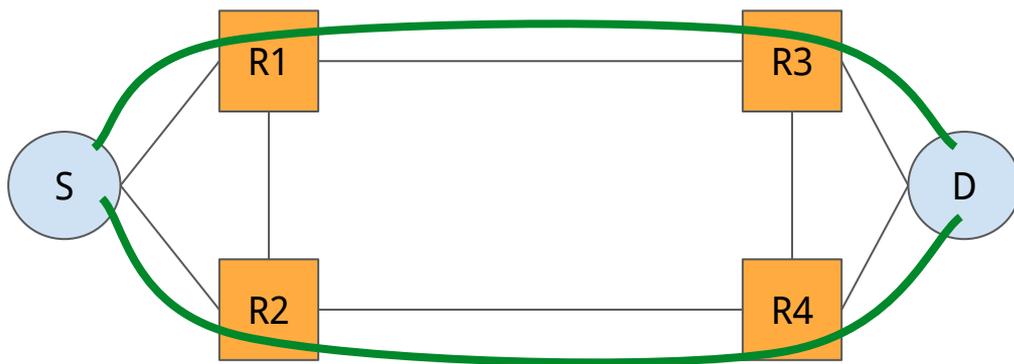
Spring 2024  
[cs168.io](https://cs168.io)

Rob Shakir

# Today

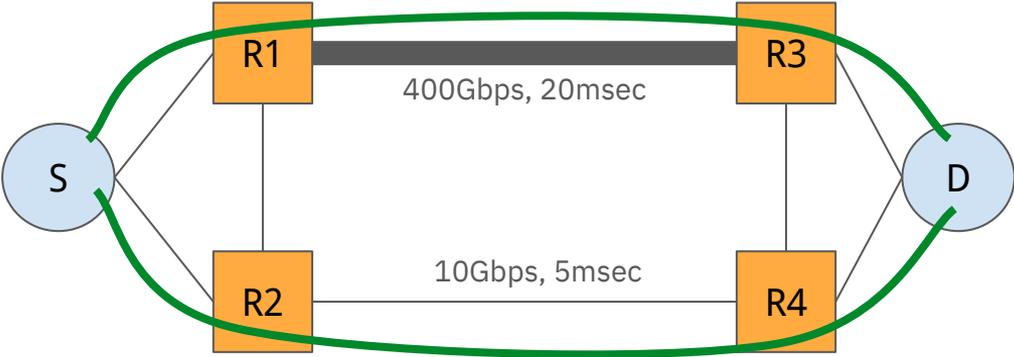
- How do we tell our networks what to do? The *management plane*.
- What do we do when “standard” networking doesn’t meet our needs?
- Software-Defined Networking.
  - In the datacenter.
  - In the wide area network.

## Recall: Least-Cost with ECMP



All costs = 1, both paths are equal.

# Recall: Least-Cost with ECMP

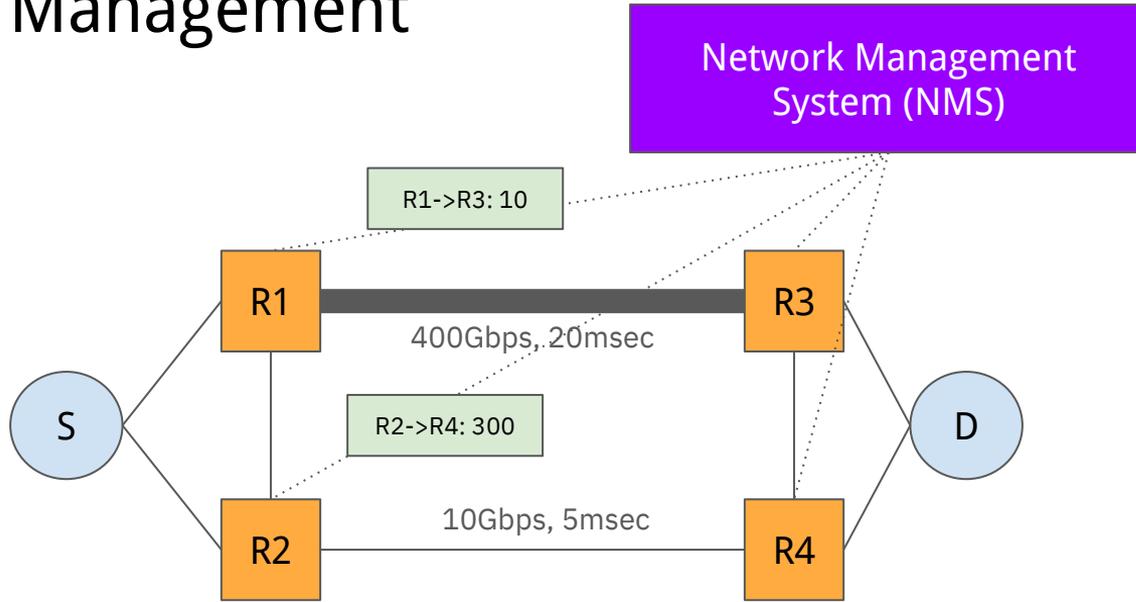


Need a way to tell a network what the costs should be.  
Recall: we said cost is *arbitrary*.

# The Management Plane

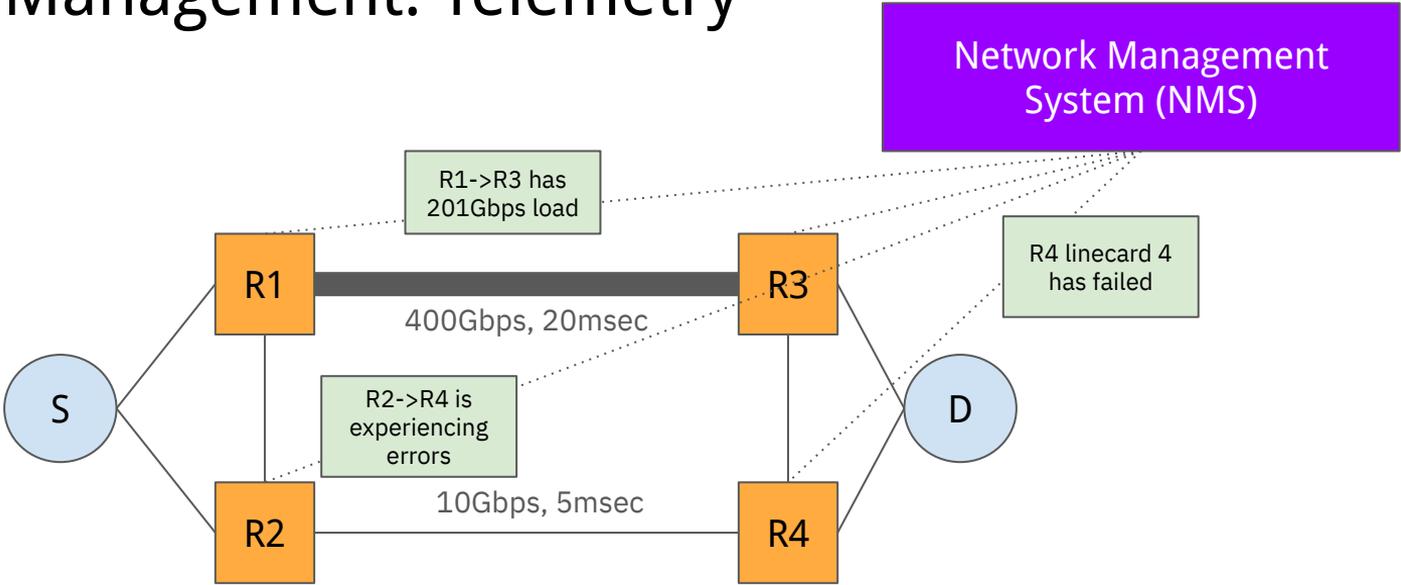
- Recall:
  - Data Plane - packet forwarding -  $O(\text{nanoseconds})$
  - Control Plane - real-time, populates device state on network devices –  $O(1 \text{ second})$
- Management plane.
  - Configuration of network device functions (including routing protocols).
  - Monitoring – statistics, alarms etc. required to run the network.
- The management plane is generally how operators tell routers what to do, and see what they are doing.
- $O(10\text{s to } 100\text{s of seconds})$

# Network Management



A network management system allows us to generate and configure parameters such as link costs.

# Network Management: Telemetry



A network management system also allows statistics and events to be read from routers.

Questions?

# How has the management plane evolved?

- There has been less focus on the management plane than other areas of networking.
  - Even though it is critical for network operation!
- Slow evolution towards using scripts to be able to programmatically control network.
  - Allowed automation of processes (e.g., adding routers and links)
  - Started to allow automation of repair of the network when things went wrong.
- But, it was the bottleneck for many network operations.

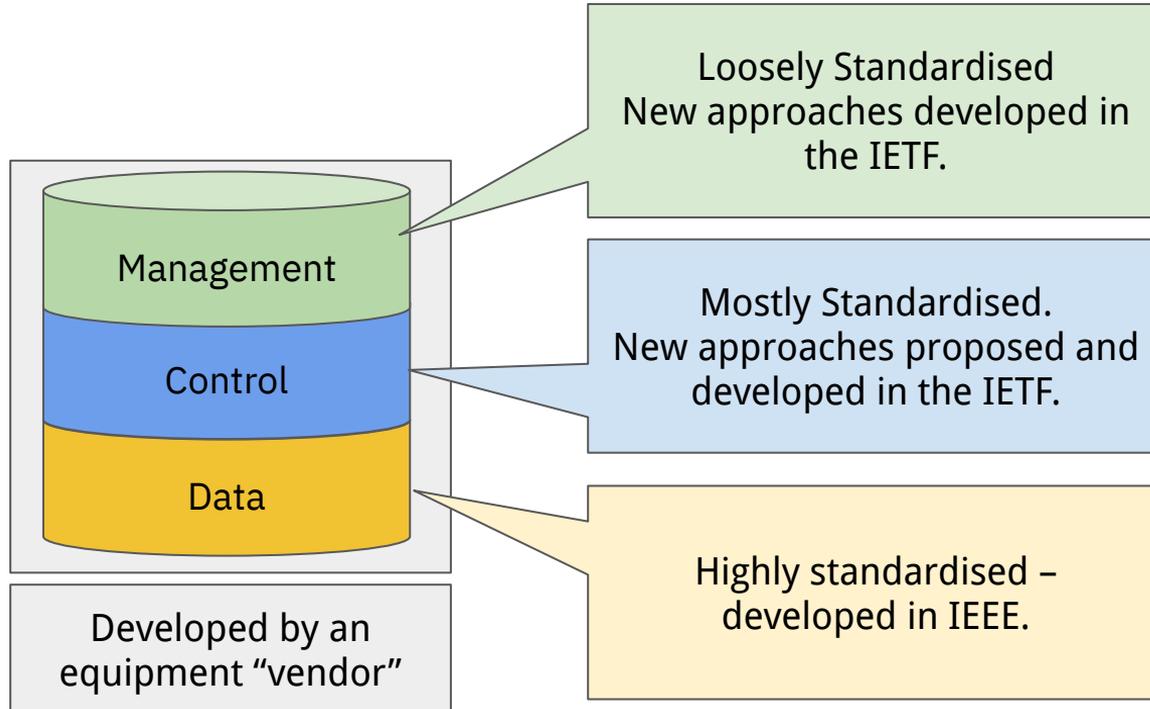
# Re-thinking the Management Plane

“Today’s data networks are surprisingly fragile and difficult to manage. We argue that the root of these problems lies in the complexity of the control and management planes”

*Greenberg et al. (2005), [A Clean Slate 4D Approach to Network Control and Management](#)*

- Researchers and network operators began to think about how better to control the network.
- This led to some more radical thinking about to redesign routers!

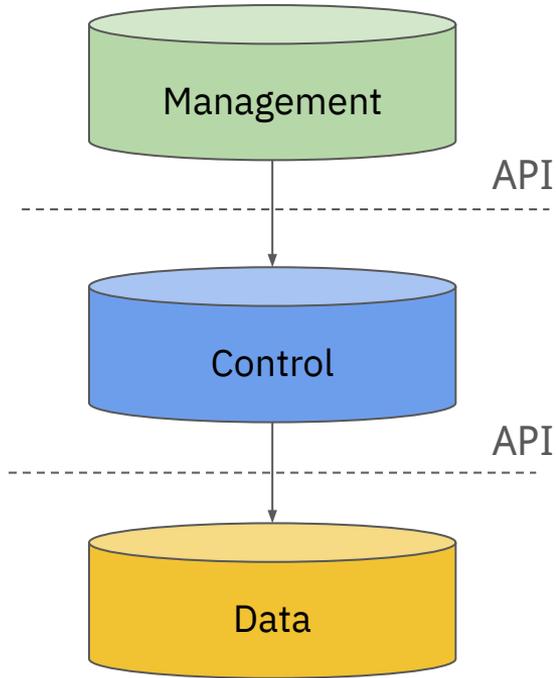
# “Traditional” Network Development Process



# Challenges of Vertical Integration

- Slow – must go through the standards process to add to the network.
- Inflexible – vendors aiming to implement solutions that work for multiple network operators.
  - What happens if my problem is different to other operators?
- Lack of ability to experiment.
  - How do I try new things in the network that I'm not sure will work out?

# Disaggregating Routers.



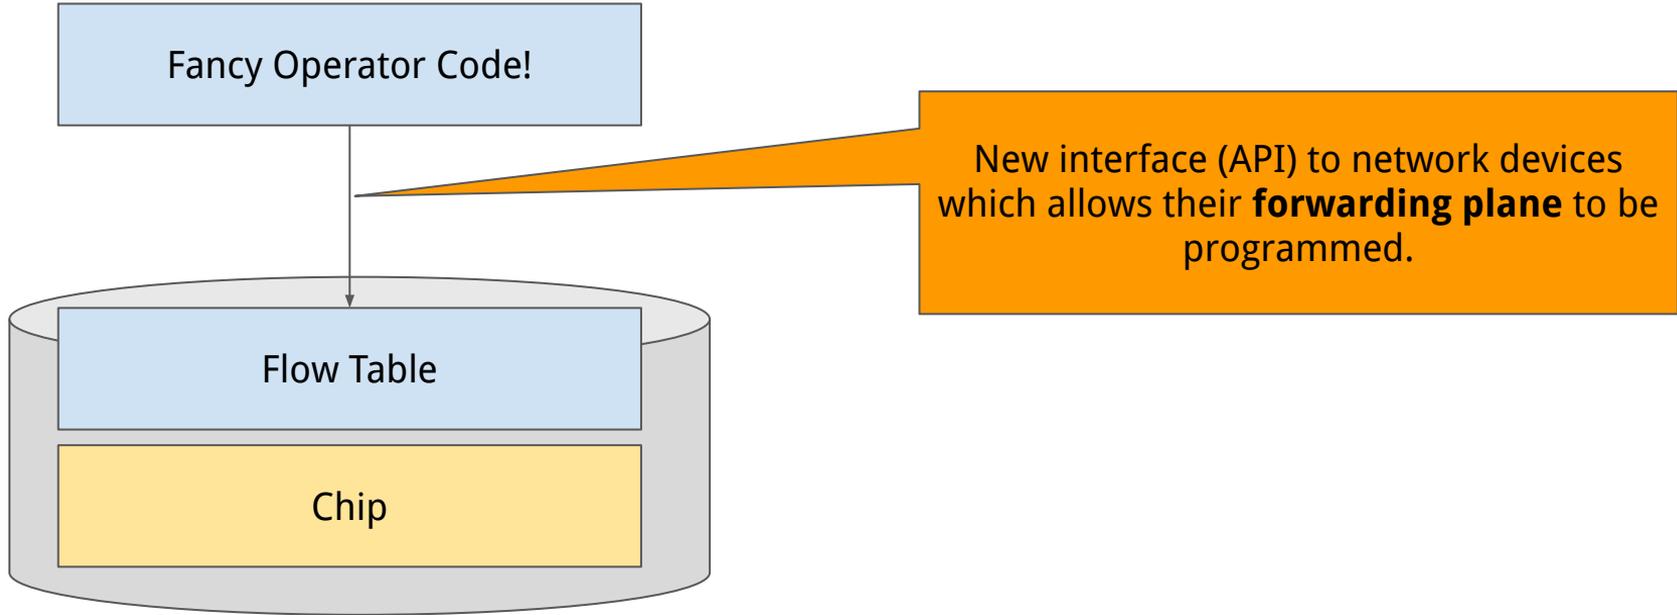
- Rather than procure everything from one place – rather split apart the network planes.
- Each plane has an API that communicates to the one below it.
  - e.g., the control plane computes forwarding tables and installs them in the data plane.
- Idea: can we allow more customisation in the control-plane by splitting the router apart?

Questions?

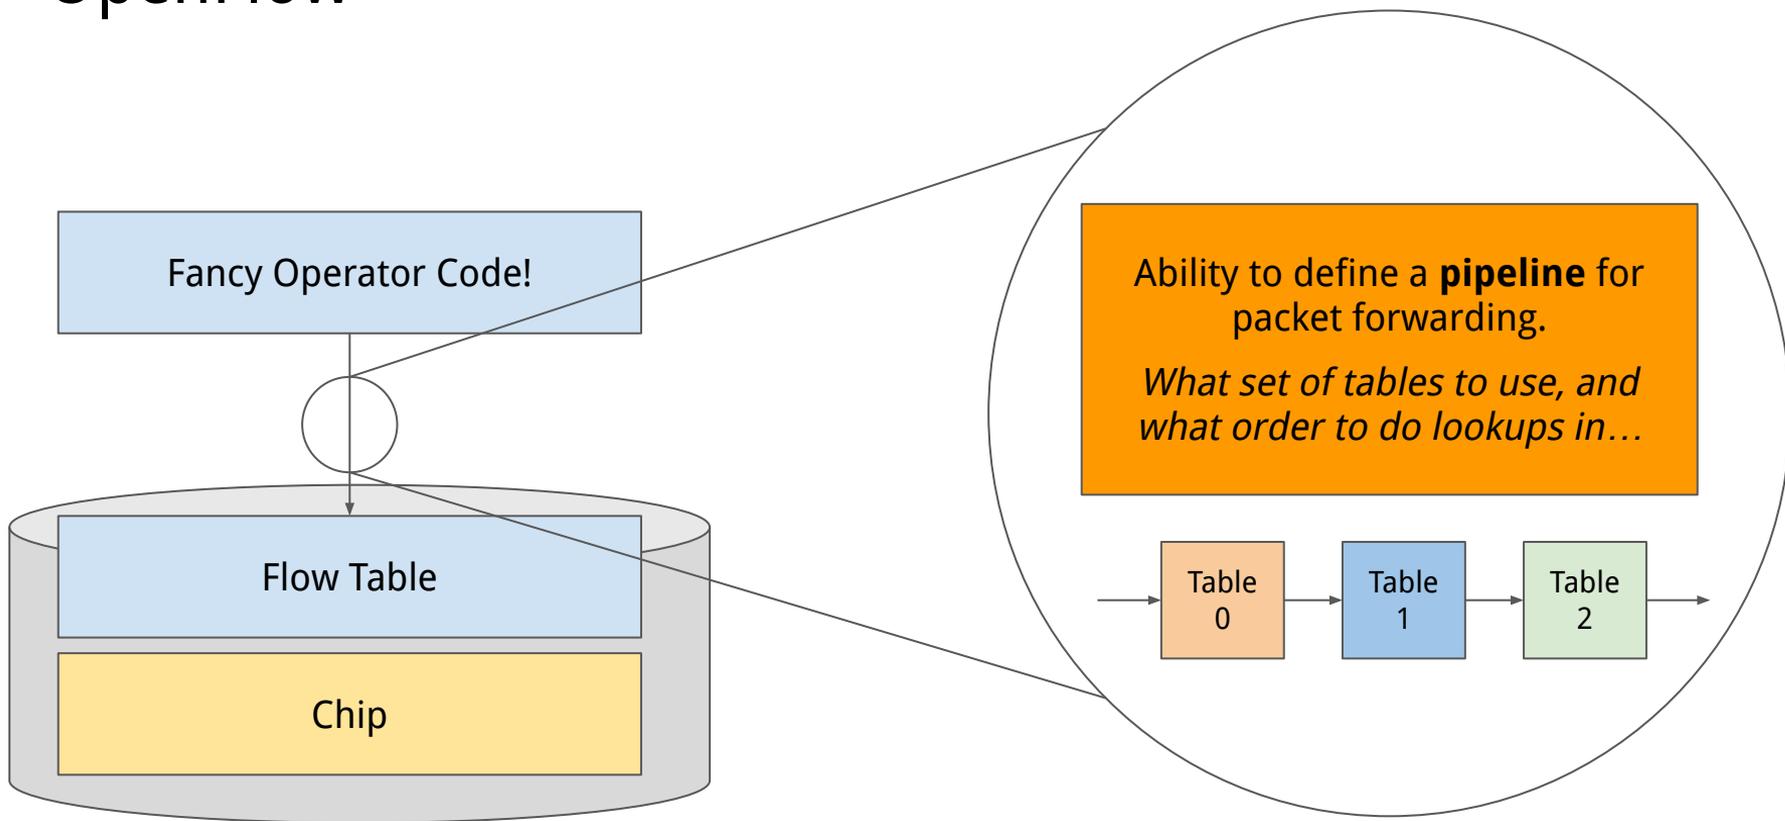
# Separating Control and Data Planes

- Not a new idea!
- IETF standardisation process examining this concept since 2003!
  - Forwarding and Control Element Separation – forces.
  - Didn't get significant momentum.
- ~2004: New management paradigms being researched.
  - RCP, 4D [Princeton, CMU]
  - **SANE, Ethane [Stanford/Berkeley]** (Scott Shenker)
- 2008: More momentum!
  - NOX Network Operating System [Nicira]
  - **OpenFlow switch interface [Stanford/Nicira]**
- 2011: Open Networking Foundation (ONF)
  - Google, Yahoo, Verizon, Deutsche Telekom, Microsoft, Facebook, NTT...
  - Cisco, Juniper, HP, Dell, Broadcom, IBM...
  - Significantly more momentum!

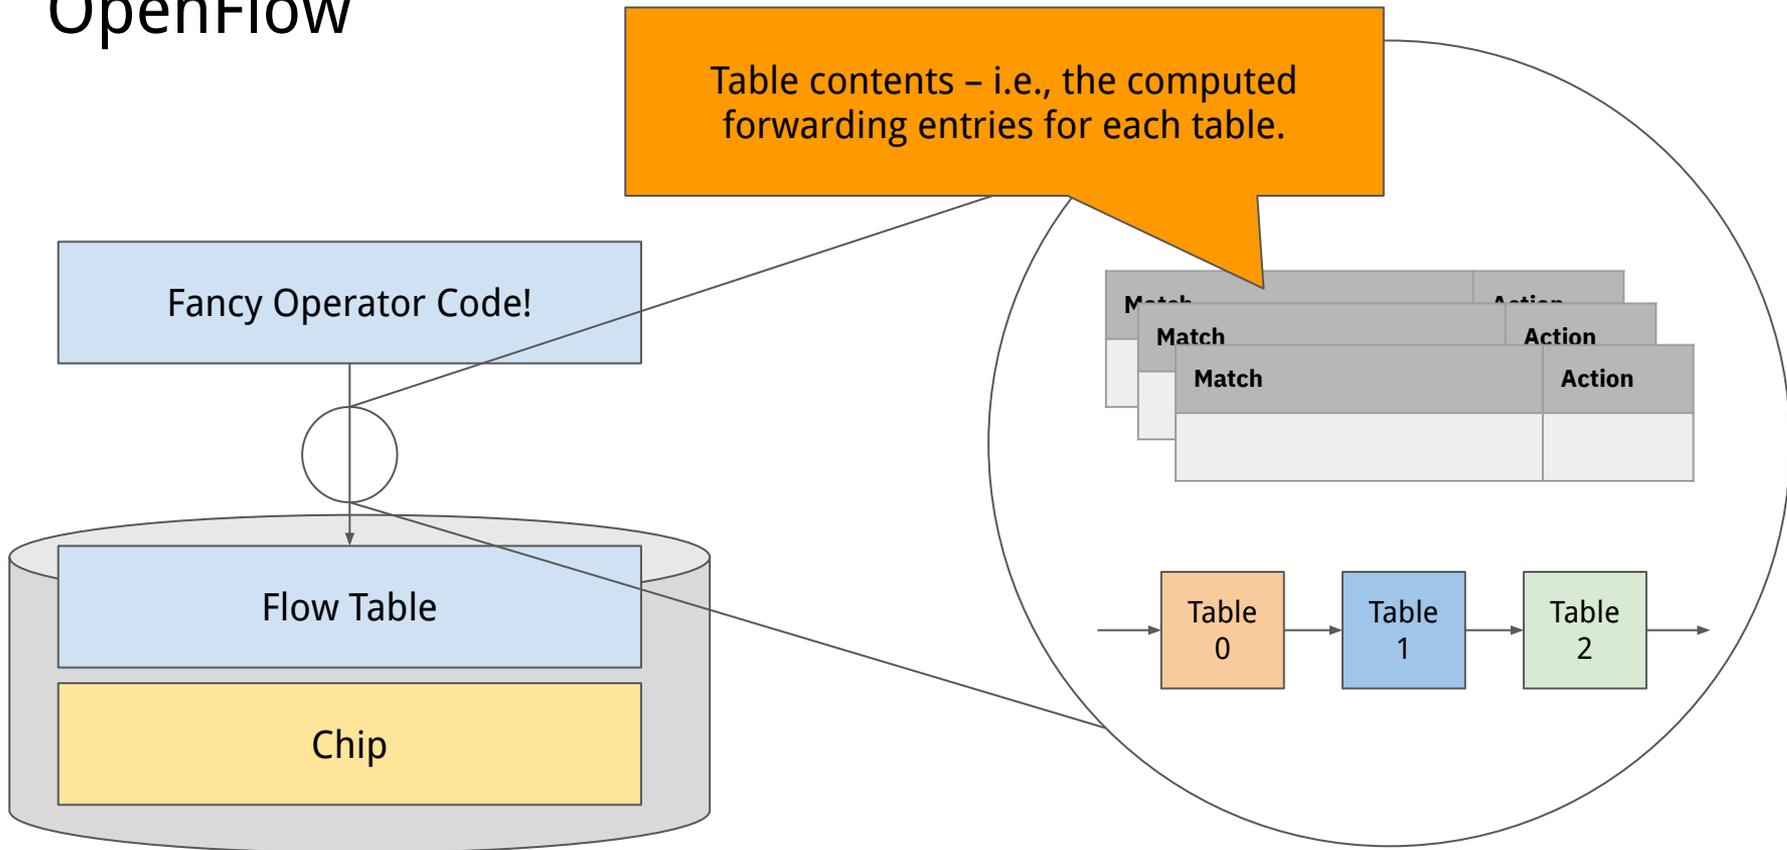
# OpenFlow



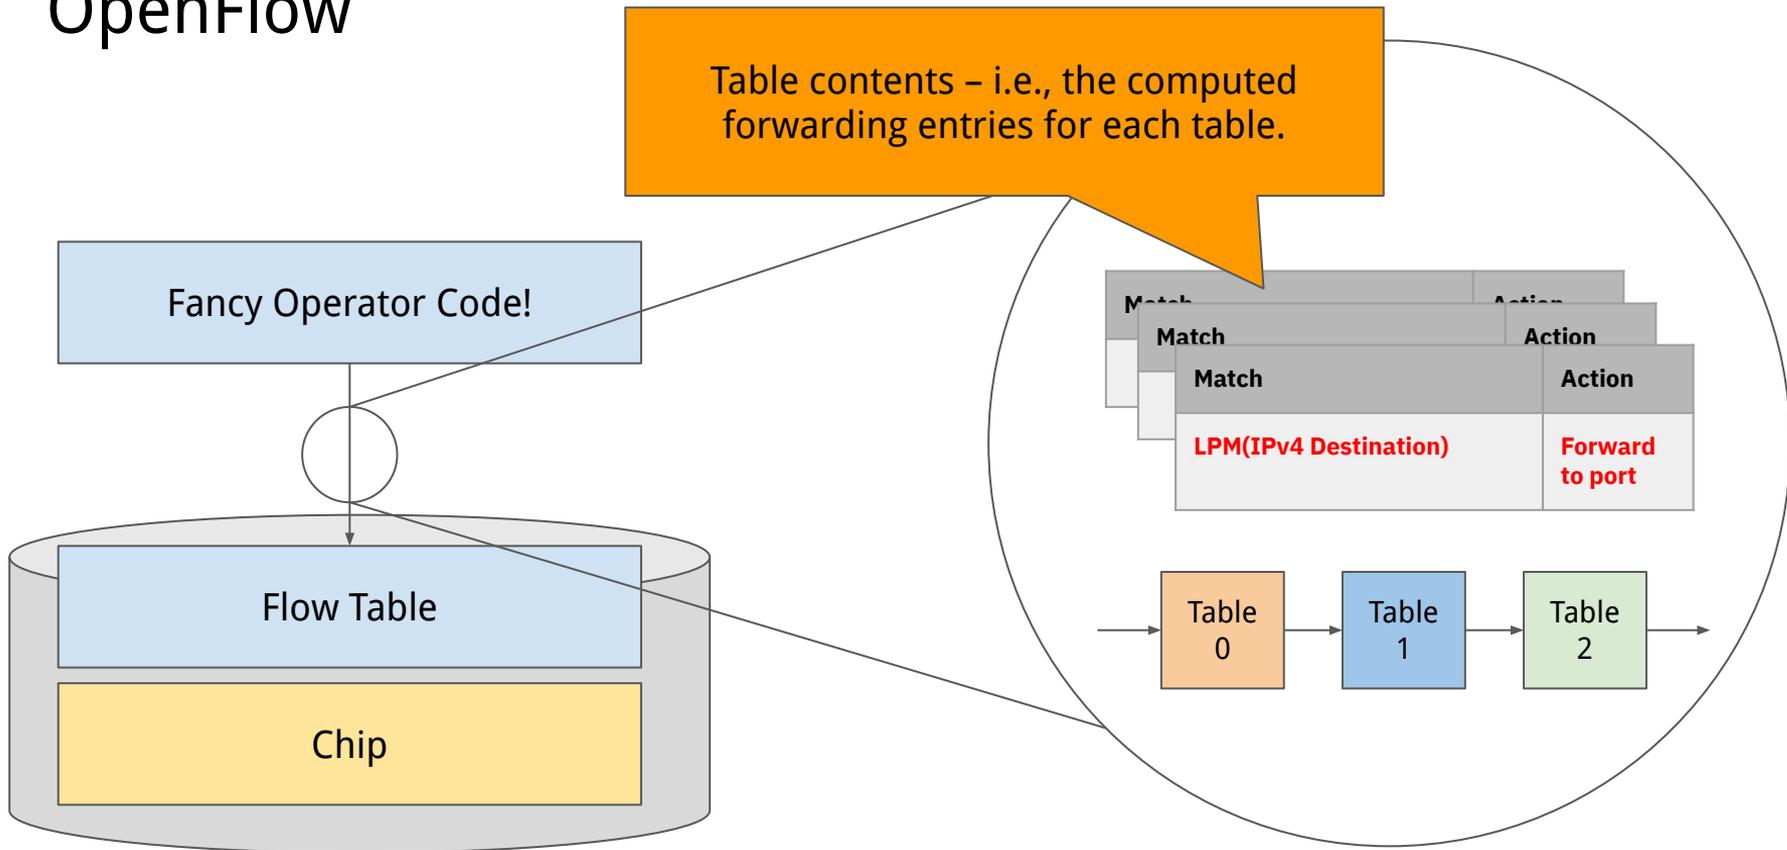
# OpenFlow



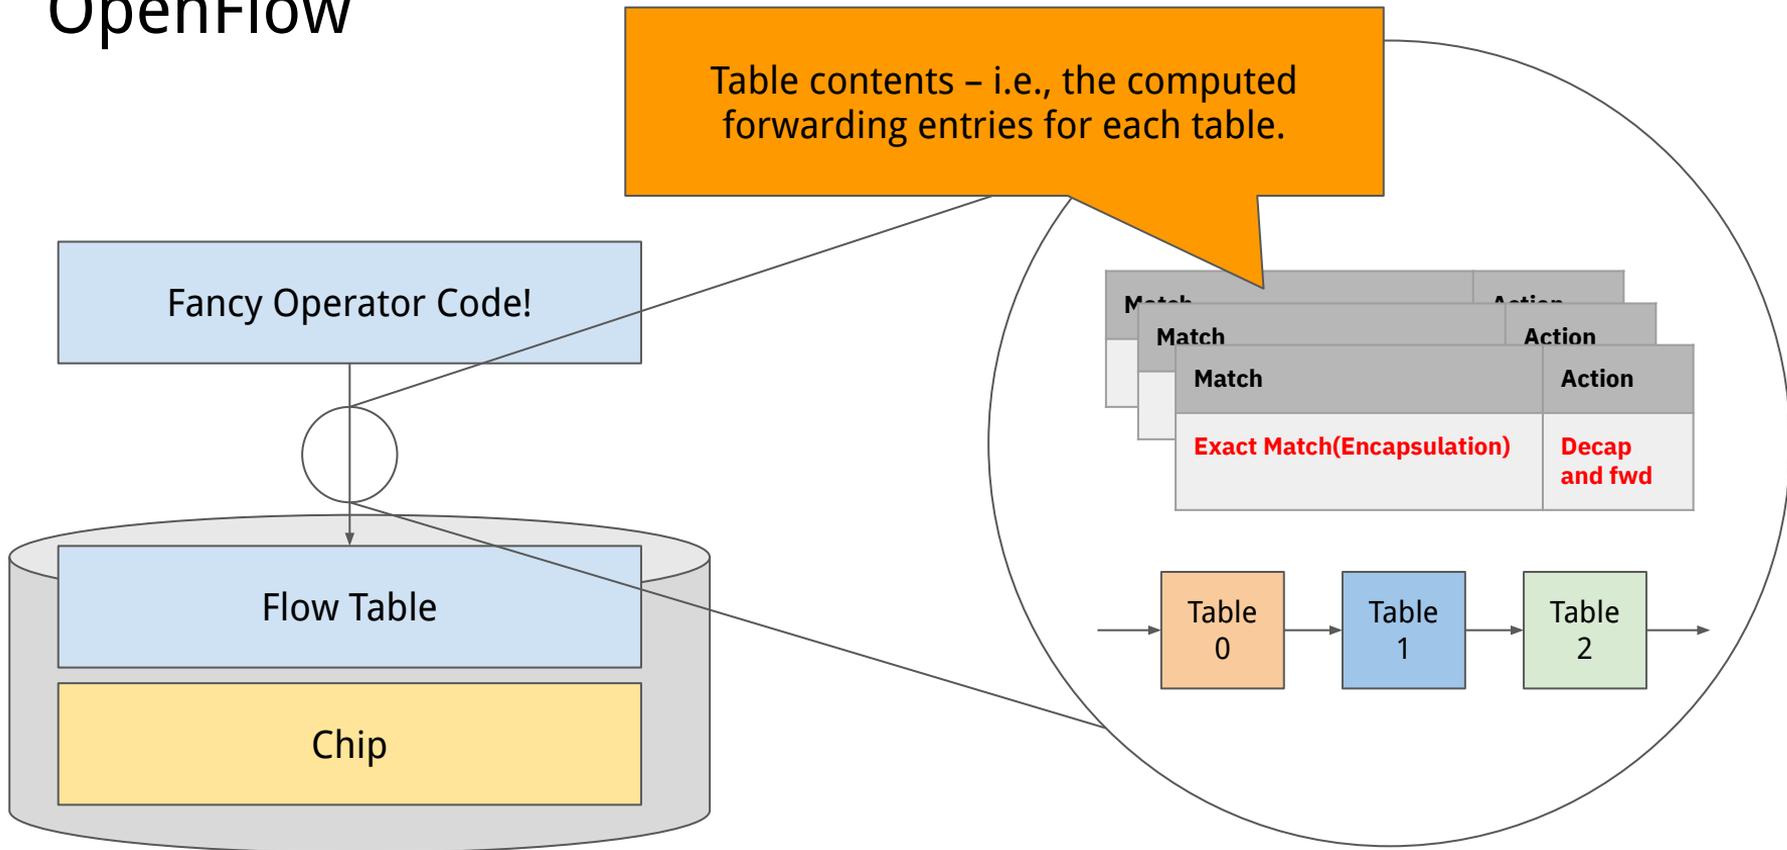
# OpenFlow



# OpenFlow



# OpenFlow



Questions?

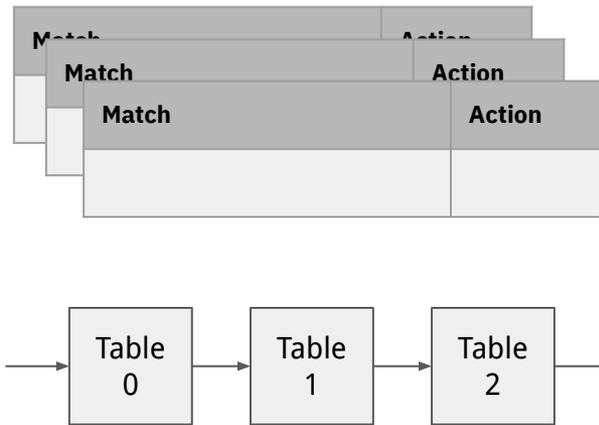
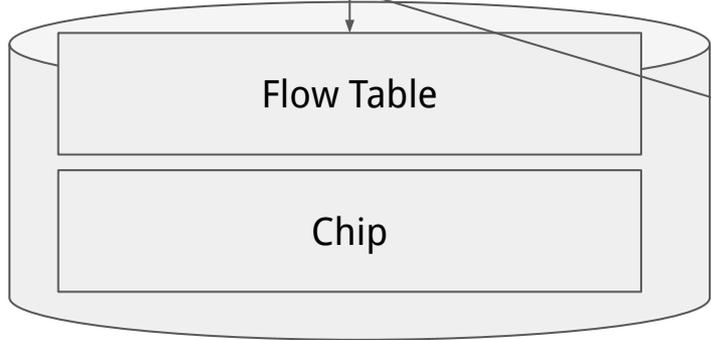
# OpenFlow Forwarding Tables

- Still constrained by what forwarding chips can do.
- Tend not to be that different to our “classic” tables.
  - IPv4 & IPv6 destination-based (LPM).
  - Policy-based routing with 5-tuples.
  - Exact match (e.g., MPLS).
- The advantage is in the flexibility of **control-plane** that this allowed.

# OpenFlow

OpenFlow was a key enabler for the control-plane to become operator specific.

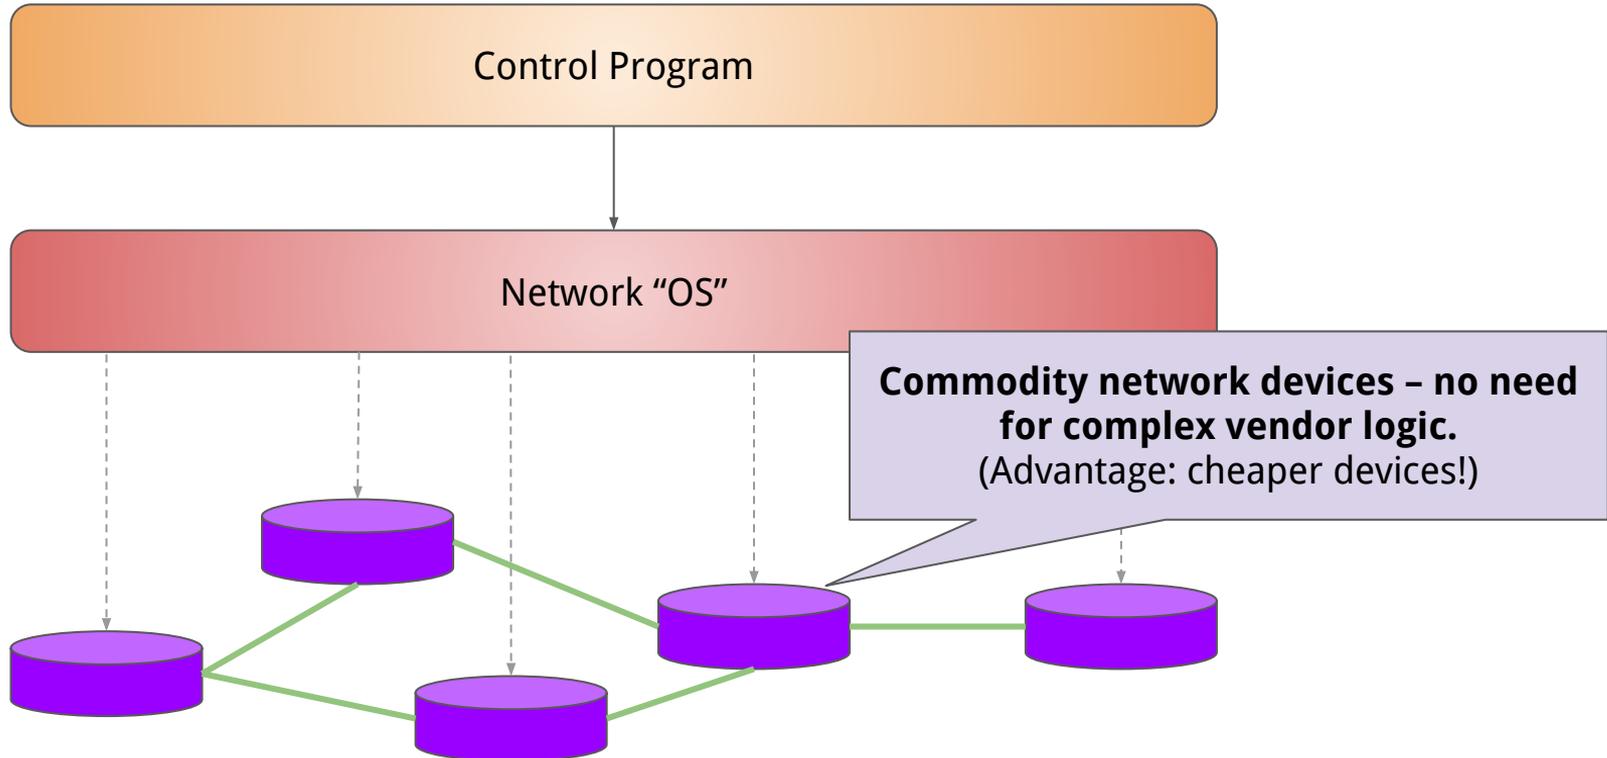
**Fancy Operator Code!**



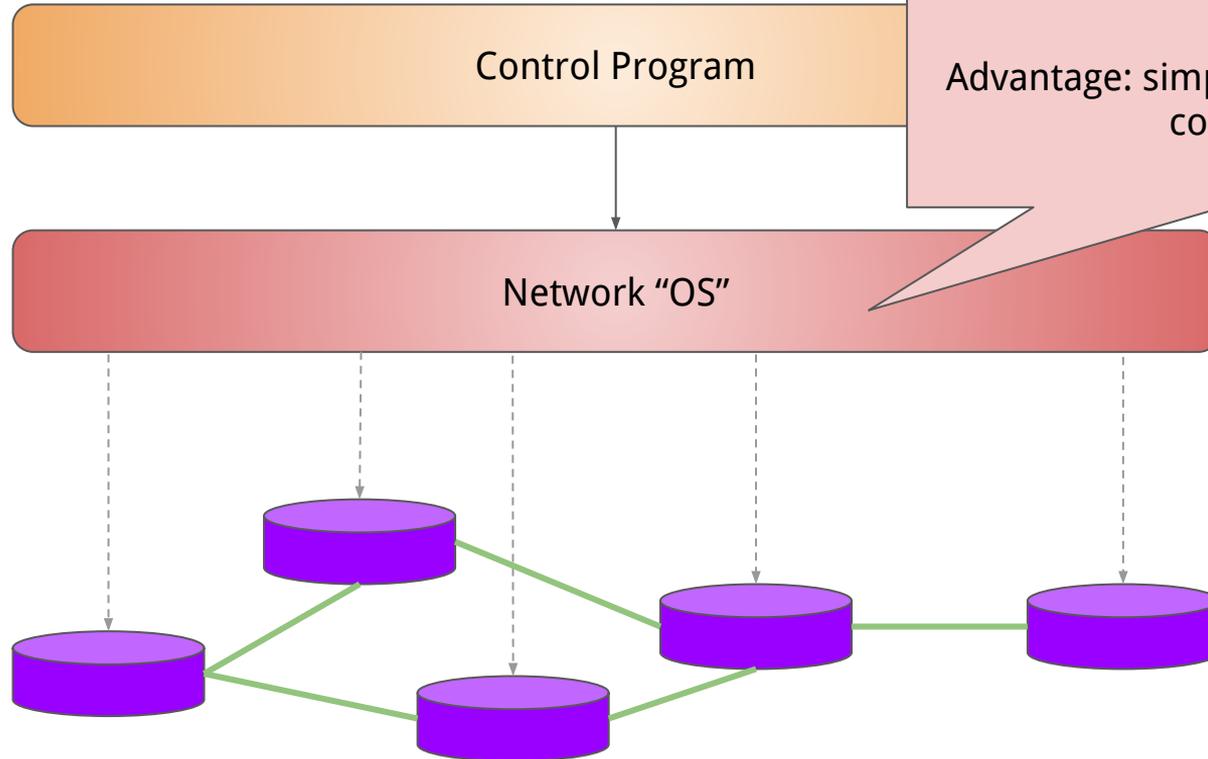
# SDN Control Planes

- An opportunity to simplify.
- The network control-plane is complex.
  - And was (and still is!) getting more and more complex based on different applications.
- Thinking about the network from a centralised viewpoint has advantages.
  - Can avoid the challenges of convergence.
  - Provides a way to make globally optimal decisions.

# The “classic” SDN view



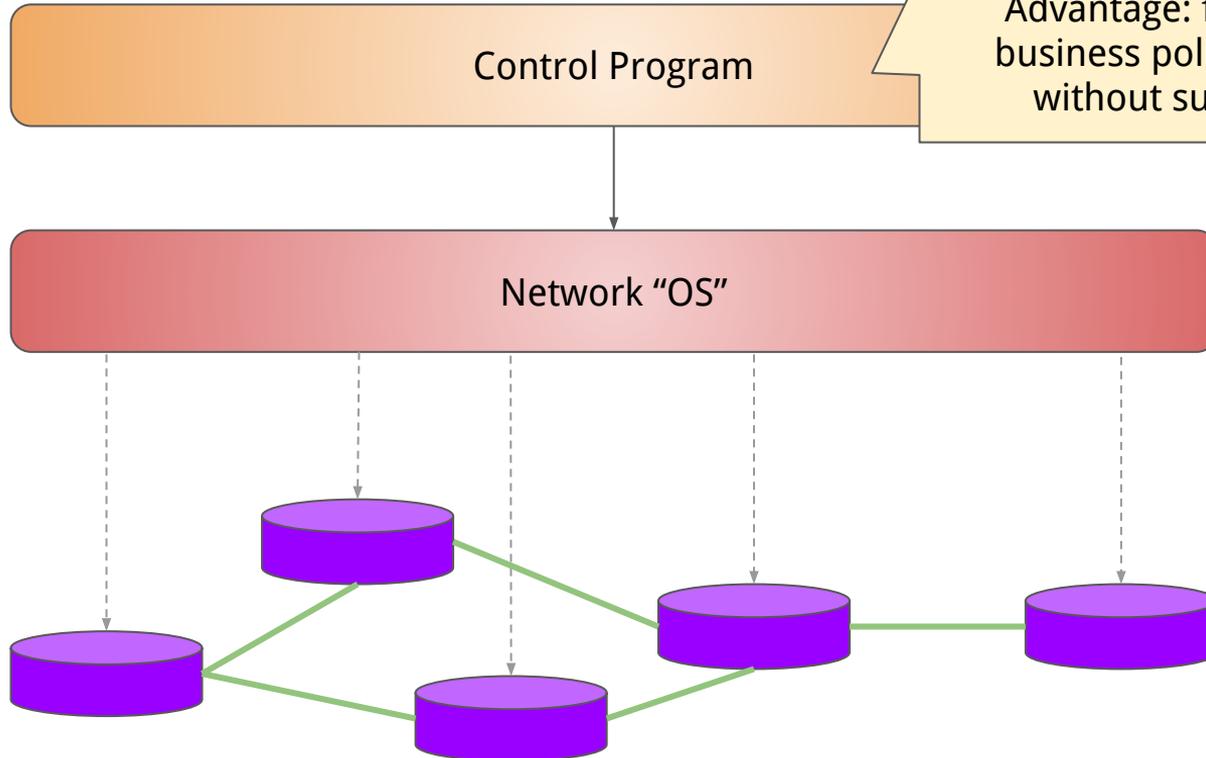
# The "classic" SDN view



**Logically centralised controller.**  
Able to provide an abstraction of the network graph and programming of the network.

Advantage: simplification of the network control plane.

# The “classic” SDN view



## **Operator-specific control programs.**

Able to consider the abstract graph of the network and make routing decisions.

Advantage: flexible means to apply business policies to network routing without supplier dependencies.

Questions?

# SDN Adoption

ACM SIGCOMM 2013

## **B4: Experience with a Globally-Deployed Software Defined WAN**

Sushant Jain, Alok Kumar, Subhasree Mandal, Joon Ong, Leon Poutievski, Arjun Singh,  
Subbaiah Venkata, Jim Wanderer, Junlan Zhou, Min Zhu, Jonathan Zolla,  
Urs Hölzle, Stephen Stuart and Amin Vahdat  
Google, Inc.  
b4-sigcomm@google.com

These characteristics led to a Software Defined Networking architecture using OpenFlow to control relatively simple switches built from merchant silicon. B4's centralized traffic engineering service drives links to near 100% utilization, while splitting application flows among multiple paths to balance capacity against application priority/demands. We describe experience with three years of B4 production deployment, lessons learned, and areas for future work.

# SDN Adoption

ACM SIGCOMM 2013

## Achieving High Utilization with Software-Driven WAN

Chi-Yao Hong (UIUC) Srikanth Kandula Ratul Mahajan Ming Zhang  
Vijay Gill Mohan Nanduri Roger Wattenhofer (ETH)

Microsoft

## Prototype

16 OpenFlow switches

- Mix of Blades and Aristas

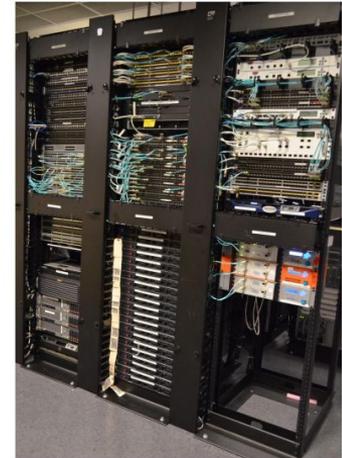
BigSwitch OpenFlow controller

32 servers as traffic sources

- 25 virtual hosts per server

8 routers (L3)

- Mix of Cisco and Juniper



# SDN Adoption

[Home](#) / Nicira OpenFlow: Networking's Next Big Thing?

## Nicira OpenFlow: Networking's Next Big Thing?

[f](#) [t](#) [in](#) [v](#) [e](#)

\$40-million startup emerges with OpenFlow platform for virtualizing the network as effectively as servers.

By **Charles Babcock**  
FEBRUARY 04, 2012

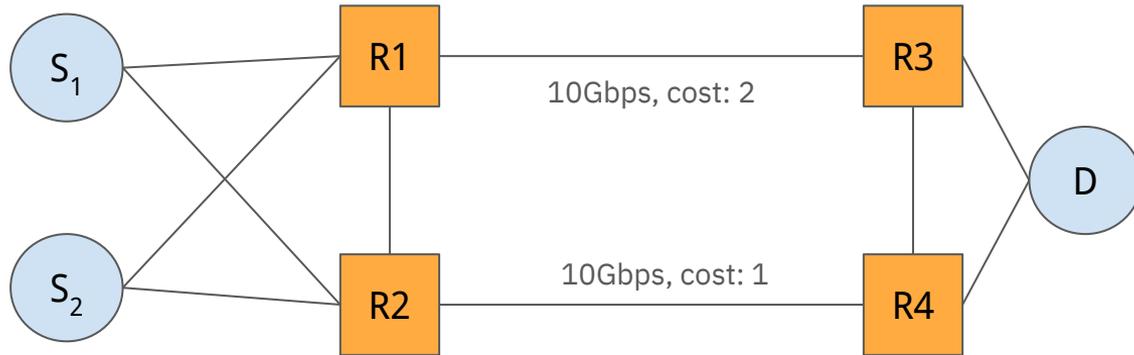
<https://www.networkcomputing.com/nicira-openflow-networkings-next-big-thing>

Not just in the wide area network – but also for datacenters  
*(we'll come back to this)*

Questions?

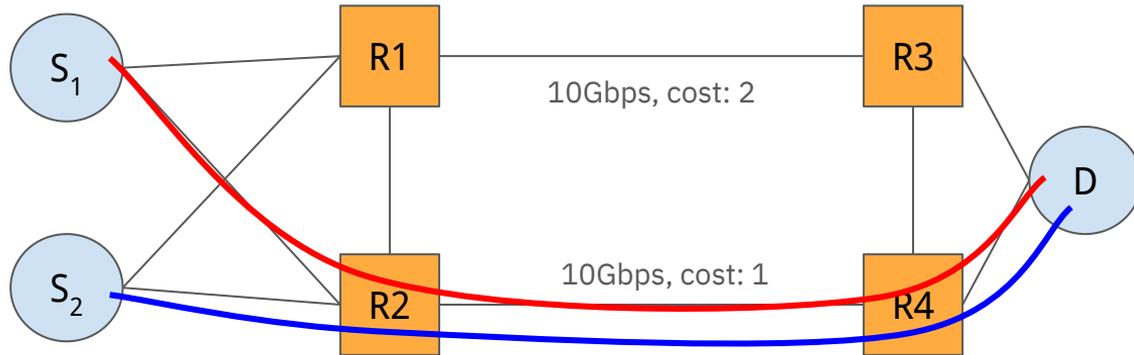
# SDN for Wide Area Networks

# Traffic Engineering



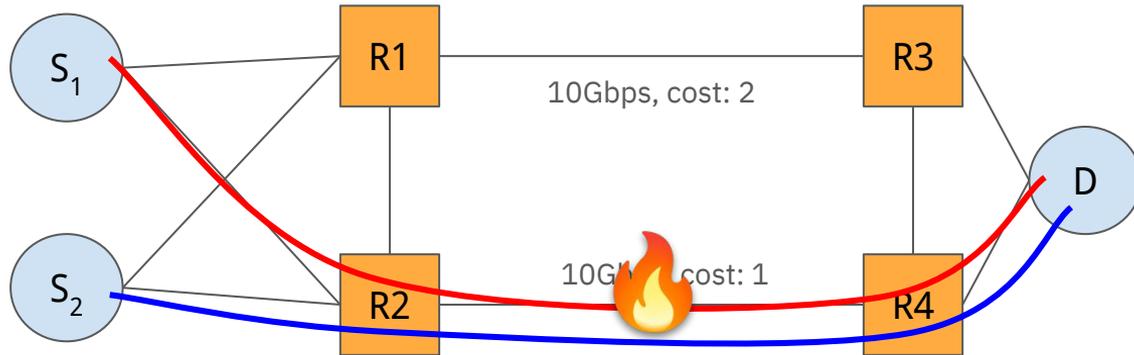
$S_1 \rightarrow D = 12\text{Gbps}$   
 $S_2 \rightarrow D = 8\text{Gbps}$

# Traffic Engineering



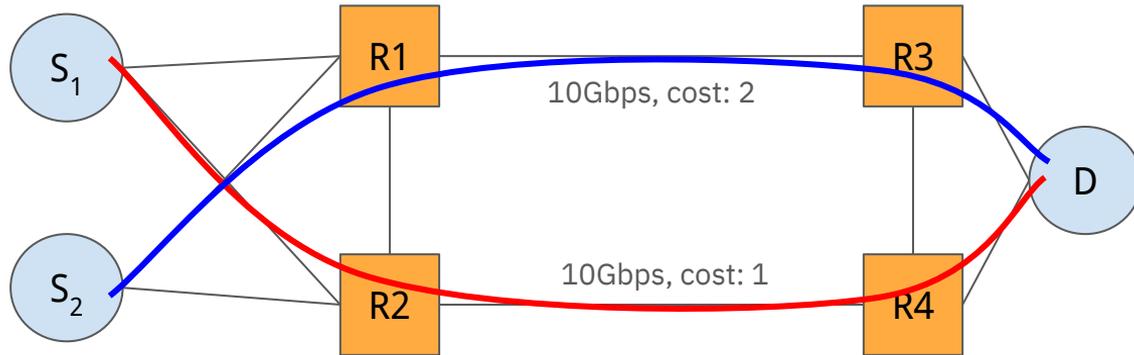
$S_1 \rightarrow D = 12\text{Gbps}$   
 $S_2 \rightarrow D = 8\text{Gbps}$

# Traffic Engineering



Route selection *without* traffic engineering results in congestion.

# Traffic Engineering

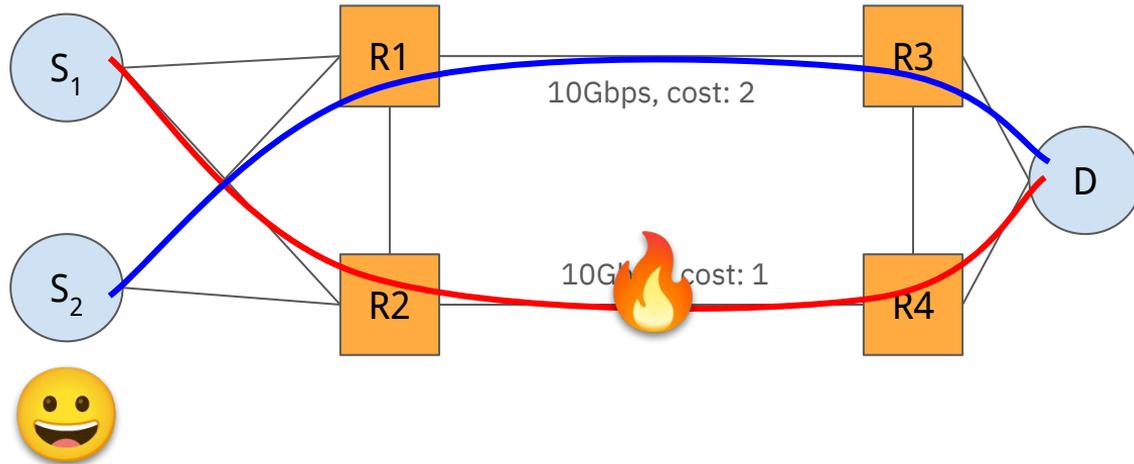


Using traffic engineering – we can tell S<sub>2</sub> to send traffic over a higher cost path.

# Traffic Engineering - cSPF

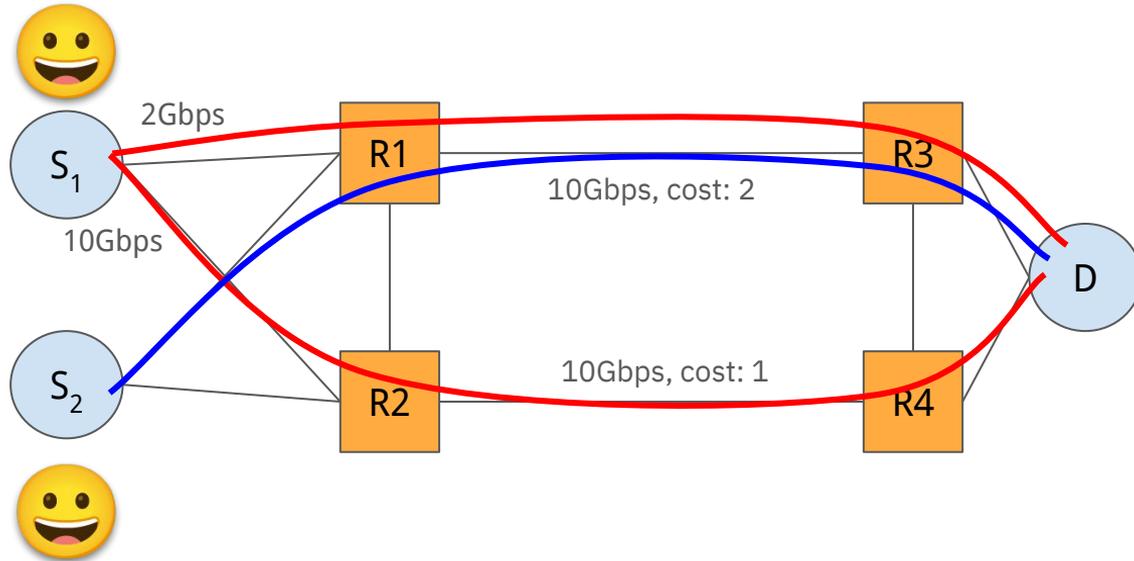
- Rather than saying “traffic should be on the shortest path”.
  - i.e., vanilla shortest-path calculation
- Say “traffic should be on the shortest path that has sufficient capacity”.
- This introduces a constraint to our shortest path calculation.
  - Hence constrained Shortest Path First.

# Traffic Engineering



Simply influencing route selection requires us to have more intelligent policy as to how to split traffic.

# Traffic Engineering

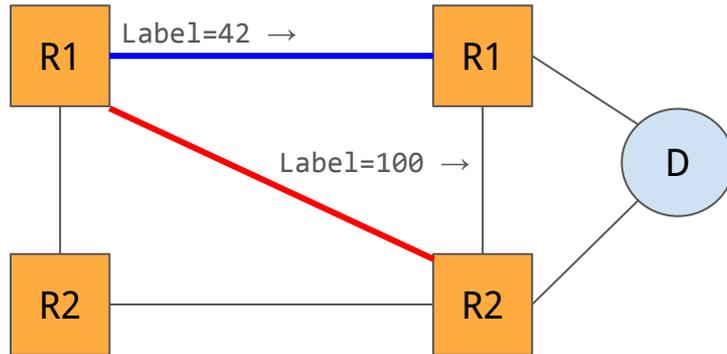
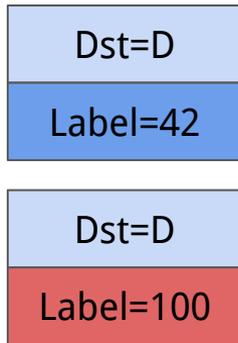


Traffic engineering policies that allow us to also split flows result in better capacity use (efficiency) of the network.

Questions?

# Traffic Engineering: Forwarding

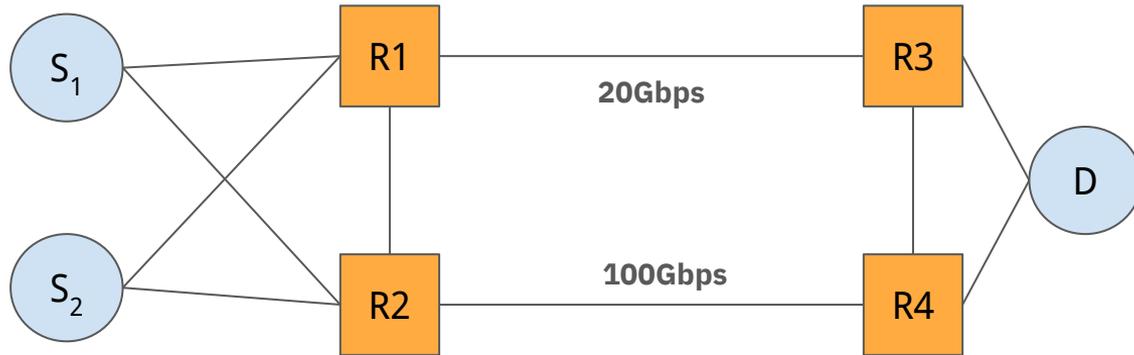
- Uses **encapsulation** like was used in the datacenter.
- Encapsulating with a new header that indicates the path to be taken through the network allows traffic engineering to be achieved with destination based forwarding.



# Centralising Traffic Engineering

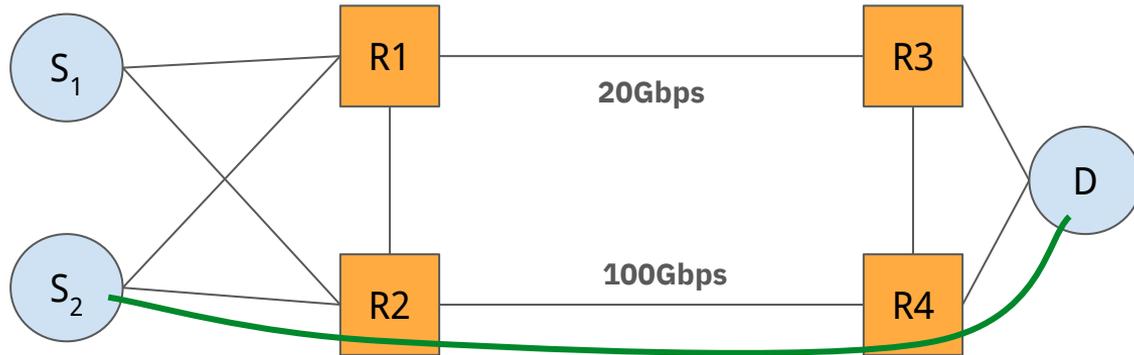
- SDN control planes allow for centralisation.
- Why might we want to centralise for TE?

# Traffic Engineering - why centralise?



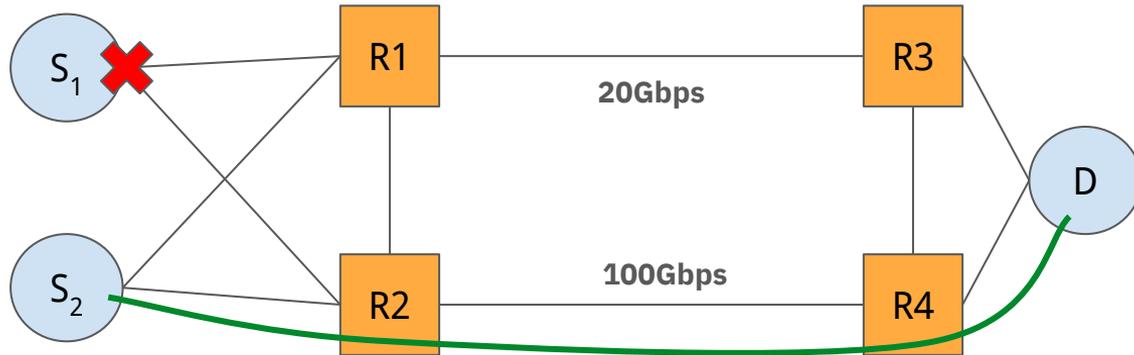
$S_1 \rightarrow D = 100\text{Gbps}$   
 $S_2 \rightarrow D = 20\text{Gbps}$

# Traffic Engineering - why centralise?



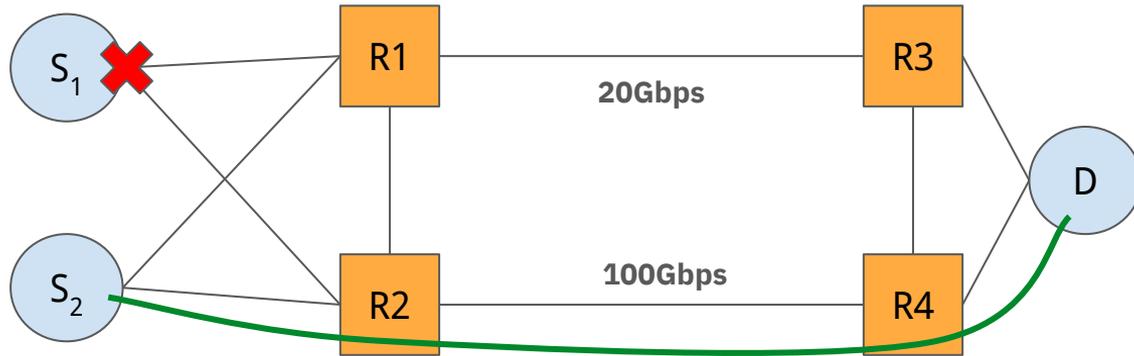
$S_2$  chooses to place its traffic via the highest bandwidth link –  
a **locally optimal** decision.

# Traffic Engineering - why centralise?



If each node acts independently, may not be able to place demands on the network.

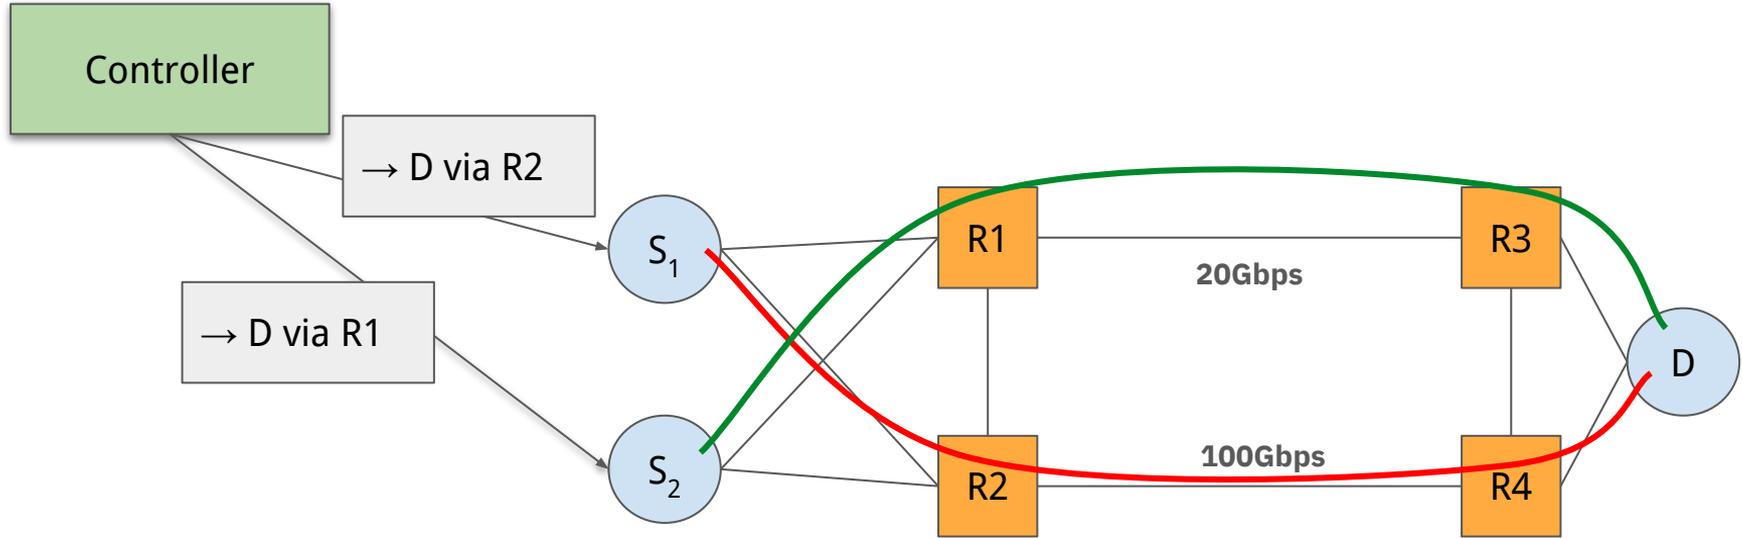
# Traffic Engineering - why centralise?



$S_1 \rightarrow D = 100\text{Gbps}$  [already split]

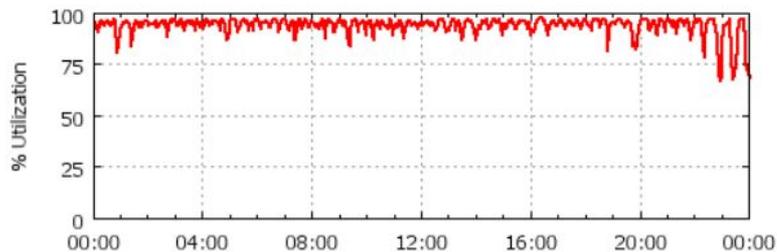
$S_2 \rightarrow D = 20\text{Gbps}$

# Traffic Engineering - why centralise?

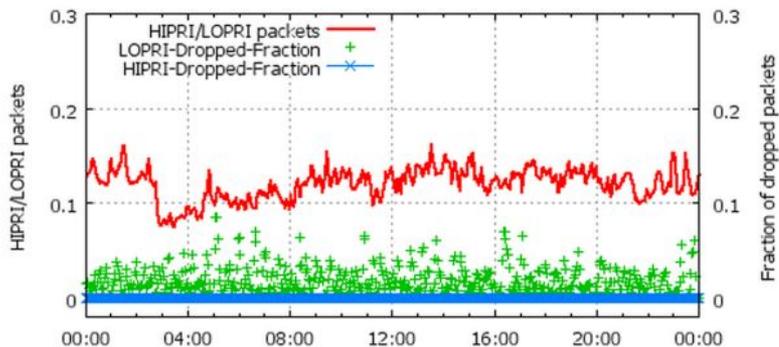


Introducing a controller allows **globally** optimal decisions to be made — increasing network efficiency.

# Utilisation Benefits



(a)



(b)

[Google B4](#)

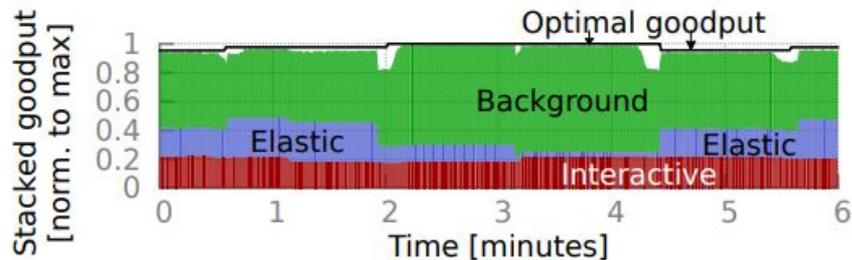


Figure 10: SWAN achieves near-optimal throughput.

[Microsoft SWAN](#)

- Significant increase in utilisation of expensive network assets (e.g., subsea fibre!).
- Ability to consider traffic that can be dropped to avoid additional build.

# Beyond Capacity Utilisation

- Off-router traffic engineering controllers can consider additional attributes and routing policies.
- Geofencing.
  - “Do not send traffic via links that are in Canada”.
- Path diversity.
  - Make  $S_1 \rightarrow D$  and  $S_2 \rightarrow D$  go via paths that never intersect.
  - Useful in applications where two paths are backups for each other – e.g., broadcast TV.

# Why might we not want to centralise?

- As with all designs – nothing comes for free.
- Reliability.
  - Traditional IP networks – one router fails, routing protocols converge around the failure.
  - SDN networks - central controller fails, network “fails static” and does not converge.
- Scalability.
  - Central controller must deal with decisions on everyone’s behalf – not just at one router.
- Complexity.
  - Infrastructure for off-device control planes etc.
- Sylvia and I are researching this right now!

Questions?

# SDN in Datacenters

## Recall: Datacenter networks

- Single owner but multiple applications hosted in them.
- Multiple tenants, with different requirements of the network.

# Private IP Addressing

- Wait? How can two hosts have the same IP address?

# Private IP Addressing

- Wait? How can two hosts have the same IP address?
- Remember, we only had  $2^{32}$  IPv4 addresses.
- Some hosts never want to be contacted from the Internet - so don't need a unique address.
- So, we can save addresses by having private addresses that can be used in multiple networks.

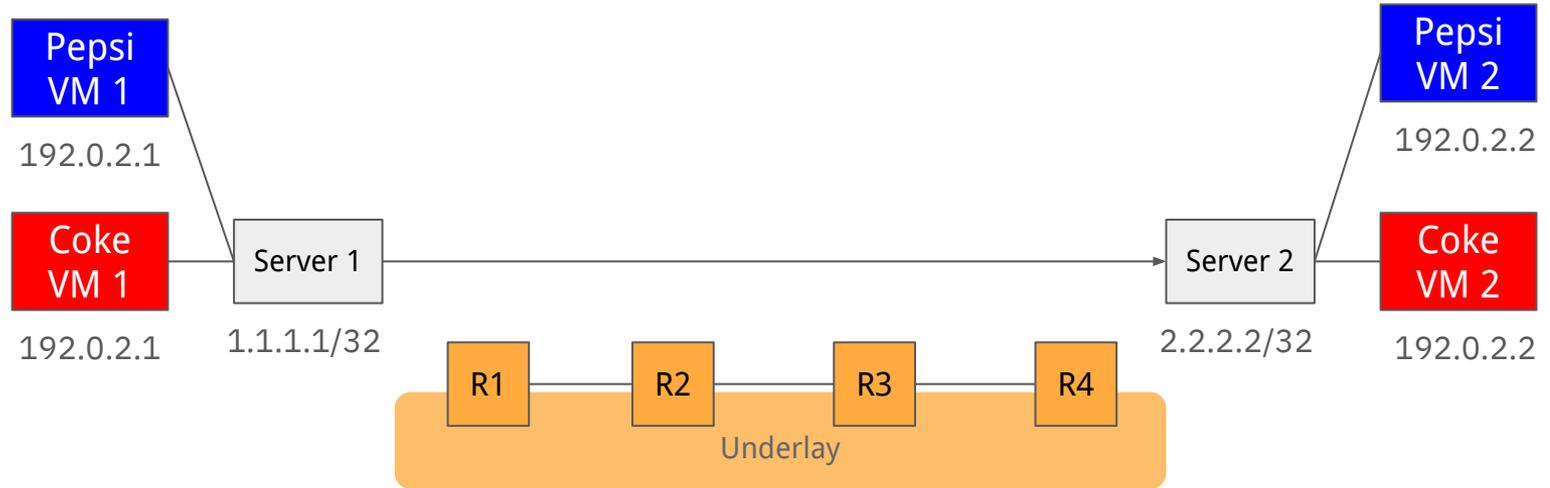
# Private IP Addressing

- Specific ranges defined for “private” use.
  - In RFC1918, so generally referred to as RFC1918 addresses.
- Most common:
  - 192.168.0.0/16
  - 10.0.0.0/8
- You’ll see these in your home network!

# Private IP Addressing

- Specific ranges defined for “private” use.
  - In RFC1918, so generally referred to as RFC1918 addresses.
- Most common:
  - 192.168.0.0/16
  - 10.0.0.0/8
- You’ll see these in your home network!
  - **But my home network talks to the Internet!**
  - We’ll come back to *Network Address Translation* (allowing multiple hosts to share the same *public* address).
    - Solution for the case that a host wants to talk *to* the Internet, but nothing makes new connections to it.

# Recall: Multi-tenancy in a Datacenter



Our datacenter networks need to support multiple tenant networks - who don't coordinate with each other.

# Per-tenant Overlay Networks

- Multiple tenants with their own *overlay networks*.
- How do we get different parts of the network to learn about the encapsulation that is needed, and where remote hosts are?

# SDN in Datacenters

192.0.2.2 Coke VM created on Server 2  
Coke Network ID is 42

SDN Controller

Table	Dst
Coke	

Coke VM 1

192.0.2.1

Server 1

1.1.1.1/32

Server 2

2.2.2.2/32

Coke VM 2

192.0.2.2

R1

R2

R3

R4

Underlay



# SDN in Datacenters

192.0.2.2 Coke VM created on Server 2  
Coke Network ID is 42

SDN Controller

Table	Dst	VNID
Coke	192.0.2.2/32	42

Coke VM 1

192.0.2.1

Server 1

1.1.1.1/32

Server 2

2.2.2.2/32

Coke VM 2

192.0.2.2

R1

R2

R3

R4

Underlay



# SDN in Datacenters

192.0.2.2 Coke VM created on Server 2  
Coke Network ID is 42

SDN Controller

Table	Dst	VNID	Remote Server
Coke	192.0.2.2/32	42	2.2.2.2/32

Coke VM 1

192.0.2.1

Server 1

1.1.1.1/32

Server 2

2.2.2.2/32

Coke VM 2

192.0.2.2

R1

R2

R3

R4

Underlay

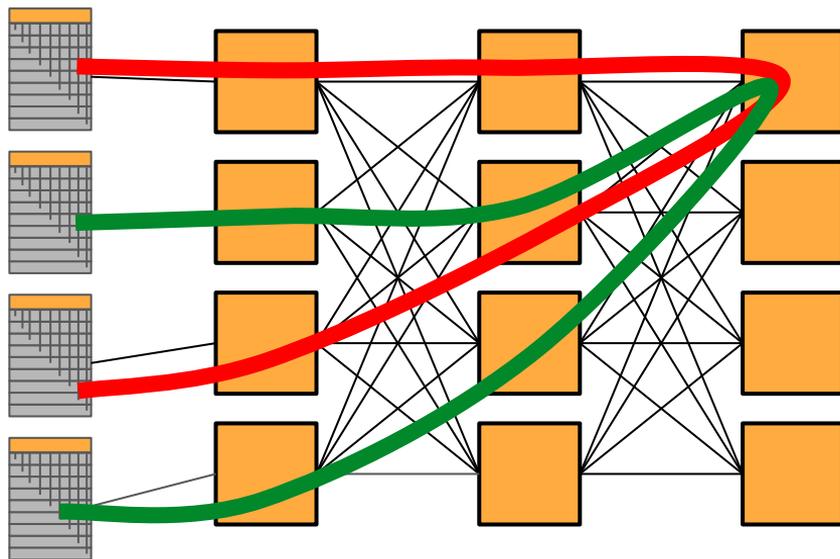
# SDN for the Overlay

- Handles distributing custom information - that would otherwise be in the control-plane of the routers.
  - e.g., virtual network IDs.
- Allows the underlay of the network to remain as simple as possible.
  - No need for the underlay to understand anything about endpoints in virtual networks.
- Allows the control plane to be extended to servers without needing routing protocols.
  - Simpler mechanism to program endpoints.

# SDN for the Datacenter underlay

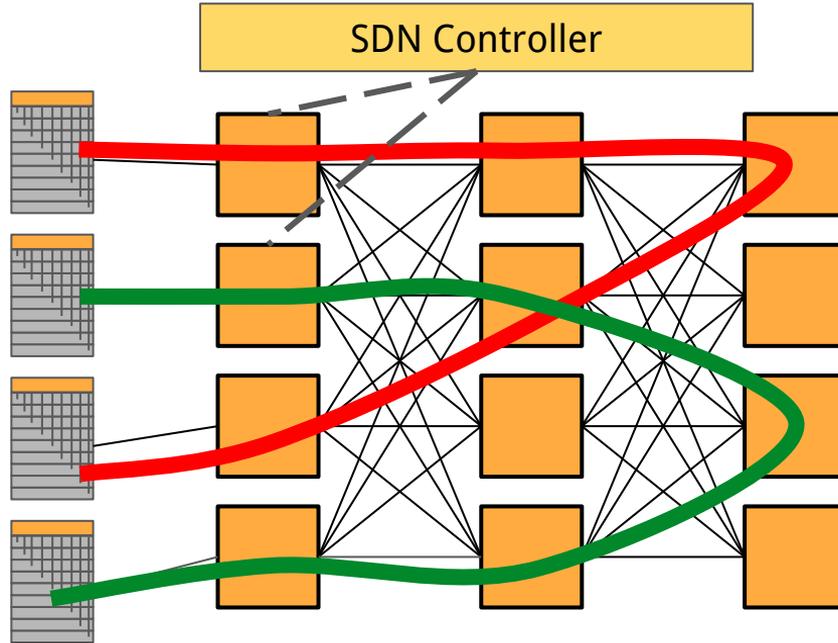
- The underlay network is just a physical network - with the same tradeoffs as designing WANs.
- Making good use of capacity, adapting to different elephant and mice flows.
- Hyperscale datacenters may use SDN in the overlay and underlay.
  - But these are decoupled systems.

# Limits of Hashing in DC Topologies



Load balancing is per-flow using the 5-tuple - elephants can still choose the same path!

# Limits of Hashing in DC Topologies



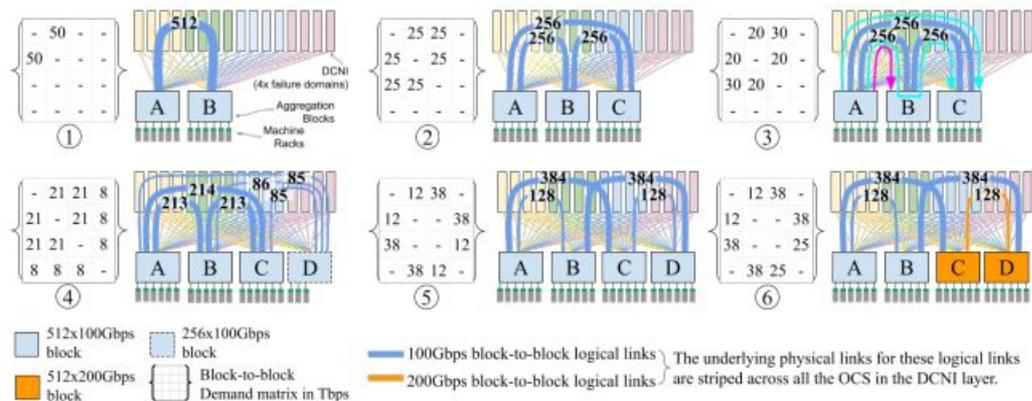
A global controller aware of different flow demands can place traffic more efficiently onto a Clos topology – improving performance.

# Improving Efficiency of Datacenters using SDN

## Jupiter Evolving: Transforming Google's Datacenter Network via Optical Circuit Switches and Software-Defined Networking

Leon Poutievski, Omid Mashayekhi, Joon Ong, Arjun Singh, Mukarram Tariq, Rui Wang, Jianan Zhang, Virginia Beauregard, Patrick Conner, Steve Gribble, Rishi Kapoor, Stephen Kratzer, Nanfang Li, Hong Liu, Karthik Nagaraj, Jason Ornstein, Samir Sawhney, Ryohei Urata, Lorenzo Vicisano, Kevin Yasumura, Shidong Zhang, Junlan Zhou, Amin Vahdat  
 Google  
 sigcomm-jupiter-evolving@google.com

- Elimination of stages in a Clos topology by adding dynamic links.
- Enabled by SDN.



Questions?

# Did SDN remove the need for the management plane?

- SDN is primarily focused on the control plane.
- But we still need some way to be able to tell routers in the network what to do, and see what they are doing.
- One of SDN's lessons is well-defined, programmatic APIs.
- This lesson also applies to the management plane of the network.

# “SDN” in the Management Plane

- Manifested itself as more software being used in networking.
- Key ideas:
  - Data modelling – YANG (IETF data modelling language)
  - OpenConfig – use of standardised APIs to interact with the management plane (like OpenFlow did for the control plane).
- Not part of the original SDN vision – but something that has evolved during its implementation.



# SDN in the Data Plane?

- Again, SDN has focused on the control-plane.
- Relatively standard set of lookup tables and behaviours available.
  - Other than in software-based dataplanes (low scale!)
- What happens if the dataplane is the inflexible part?
  - Programming Protocol-Independent Packet Processors (P4)
  - Open source, domain-specific language for telling a device how to process packets.
  - Programmability at the data plane.
- Some challenges!
  - Fixed functionality forwarding chips (needed for speed) are not arbitrarily flexible.
  - Programmable chips are lower performance, and higher cost.

Questions?

# Recap

- SDN evolved from inflexibility of the control-plane to meet different operational requirements.
- It splits the control- and data-plane – to make the control-plane more flexible.
- There are applications across both WAN and datacenter networks.