

SDN Usecases

ECE/CS598HPN

Radhika Mittal

Logistics

- Do all of you receive my emails?
- Are you all submitting your reading assignments?
- Do you have access to Illinois media space?
- Warm-up assignment due on Thursday. Have all of you found grading partners?
- Sign up for the project proposal meeting next week!
- Would you like your opinions to be anonymous or is name calling ok?

B4: Experience with a Globally- Deployed Software Defined WAN

Google

SIGCOMM'13

B4: Google's Software-Defined WAN

- Google operates two separate backbones:
 - B2: carries Internet facing traffic
 - Growing at a rate faster than the Internet
 - B4: carries inter-datacenter traffic
 - More traffic than B2
 - Growing faster than B2

B4: Google's Software-Defined WAN

- Google operates two separate backbones:
 - B2: carries Internet facing traffic
 - Growing at a rate faster than the Internet
 - **B4: carries inter-datacenter traffic**
 - More traffic than B2
 - Growing faster than B2

B4: Google's Software-Defined WAN

Among the first and largest SDN/OpenFlow deployment.



Why SDN/OpenFlow?


- Opportunity to reason about global state
 - Simplified coordination and orchestration.
- Exploit raw speed of commodity servers.
 - Latest generation servers are much faster than embedded switch processors.
- Decouple software and hardware evolution.
 - Control plane software can evolve more quickly.
 - Data plane hardware can evolve slower based on programmability and performance.

What did B4 use SDN for?

- Centralized routing.
 - Basic functionality.
 - Allowed Google to develop and stress test the SDN architecture.
- Centralized traffic engineering.
 - Allocating routes (and bandwidth) to groups of flows.
 - Also allows prioritizing some flows over others.
 - Enables running the WAN at higher utilization.

Traffic Engineering


- Traditionally accomplished via MPLS tunnels.
 - Tunnels defines routes and priority.
 - Ingress routers locally and greedily map flows to tunnels.



- Centralized TE using SDNs allows closer to optimal routes.


Traffic Engineering: another example

Flows: R1->R6: 20; R2->R6: 20; R4->R6: 20




Traffic Engineering: another example

Flows: R1->R6: 20; R2->R6: 20; R4->R6: 20



Traffic Engineering: another example

Flows: R1->R6: 20; R2->R6: 20; R4->R6: 20




R5-R6 link fails

- R1, R2, R4 *autonomously* find next best path

Traffic Engineering: another example

Flows: R1->R6: 20; R2->R6: 20; R4->R6: 20




R5-R6 link fails

- R1, R2, R4 *autonomously* try for next best path
- R1, R2, R4 push **20** altogether

No Traffic Engineering

Traffic Engineering: another example

Flows: R1->R6: 20; R2->R6: 20; R4->R6: 20



R5-R6 link fails


- R1, R2, R4 *autonomously* try for next best path
- R1 wins, R2, R4 retry for next best path

Distributed Traffic Engineering Protocols

e.g. MPLS + RSVP

Traffic Engineering: another example

Flows: R1->R6: 20; R2->R6: 20; R4->R6: 20



R5-R6 link fails


- R1, R2, R4 *autonomously* try for next best path
- R1 wins, R2, R4 retry for next best path
- R2 wins this round, R4 retries again

Distributed Traffic Engineering Protocols

e.g. MPLS + RSVP

Traffic Engineering: another example

Flows: R1->R6: 20; R2->R6: 20; R4->R6: 20



R5-R6 link fails


- R1, R2, R4 *autonomously* try for next best path
- R1 wins, R2, R4 retry for next best path
- R2 wins this round, R4 retries again
- R4 finally gets third best path!

Distributed Traffic Engineering Protocols

e.g. MPLS + RSVP

Traffic Engineering: another example

Simple topology



Flows:

- R1->R6: 20; R2->R6: 20; R4->R6: 20

R5-R6 fails

- R5 informs TE, which programs routers in one shot

Centralized Traffic Engineering Protocols

Limitation of OpenFlow faced by B4

- Needs somewhat fancier switch behavior.
 - TE enforced using IP-in-IP tunnels.
 - Switches should understand how to parse headers for tunneling.
 - Encapsulate with tunnel IP at source ingress.
 - Decapsulate tunnel IP and destination egress.
- Developed their own switches that supported a slightly extended version of OpenFlow.

B4 SDN architecture

protocol

silicon

protocol

silicon

protocol

silicon

protocol

silicon

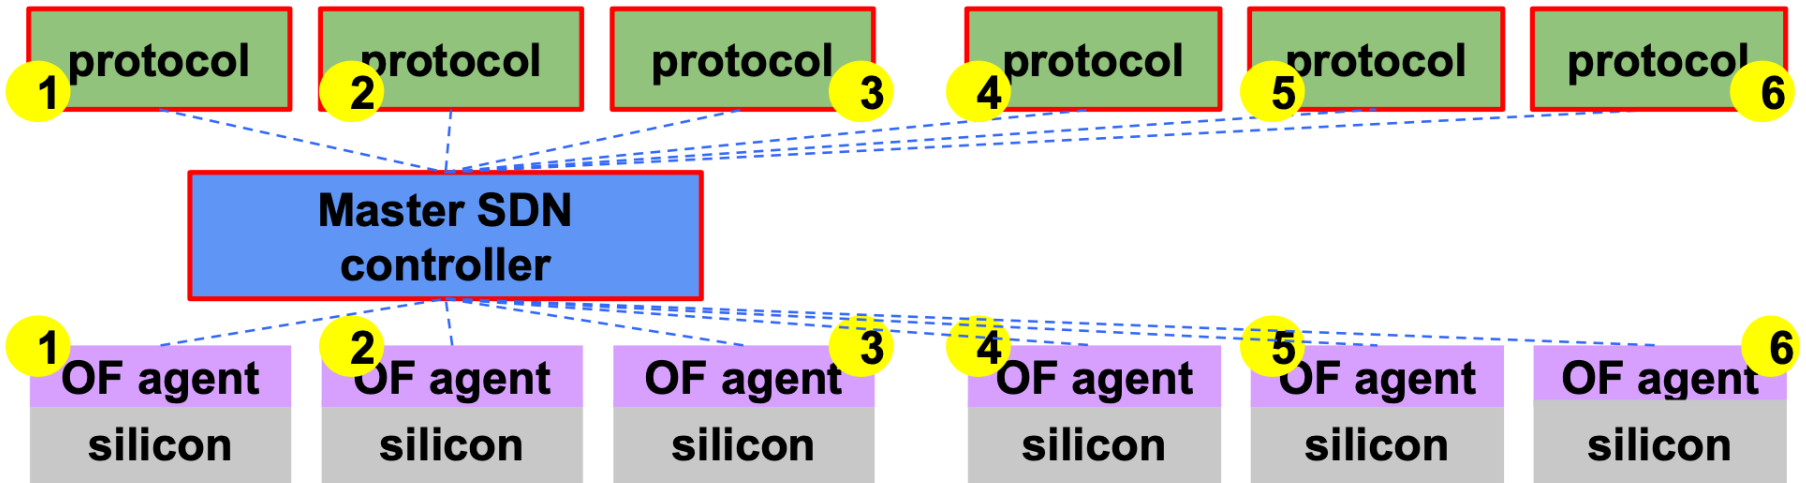
protocol

silicon


protocol

silicon

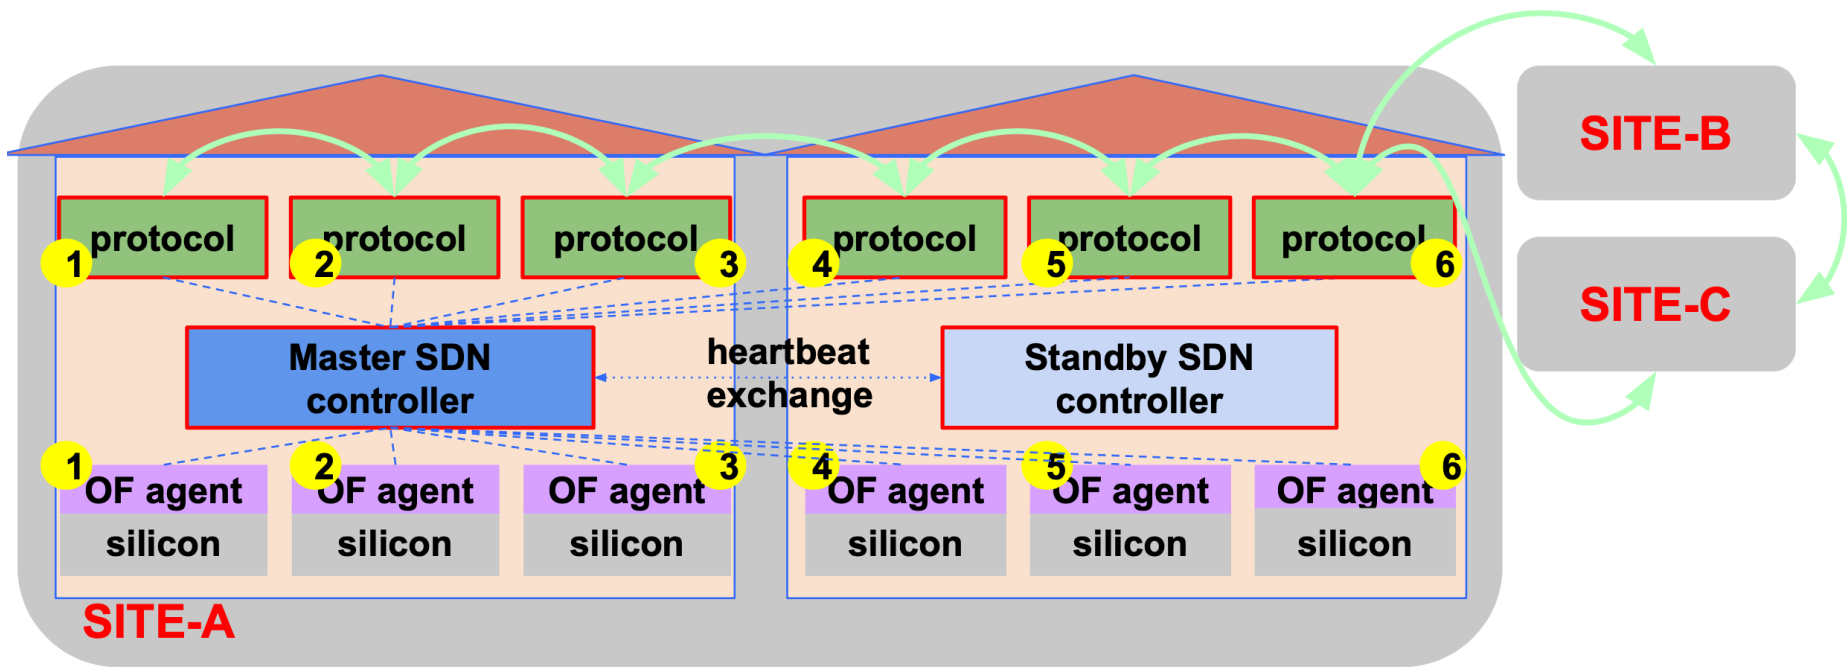
B4 SDN architecture



B4 SDN architecture

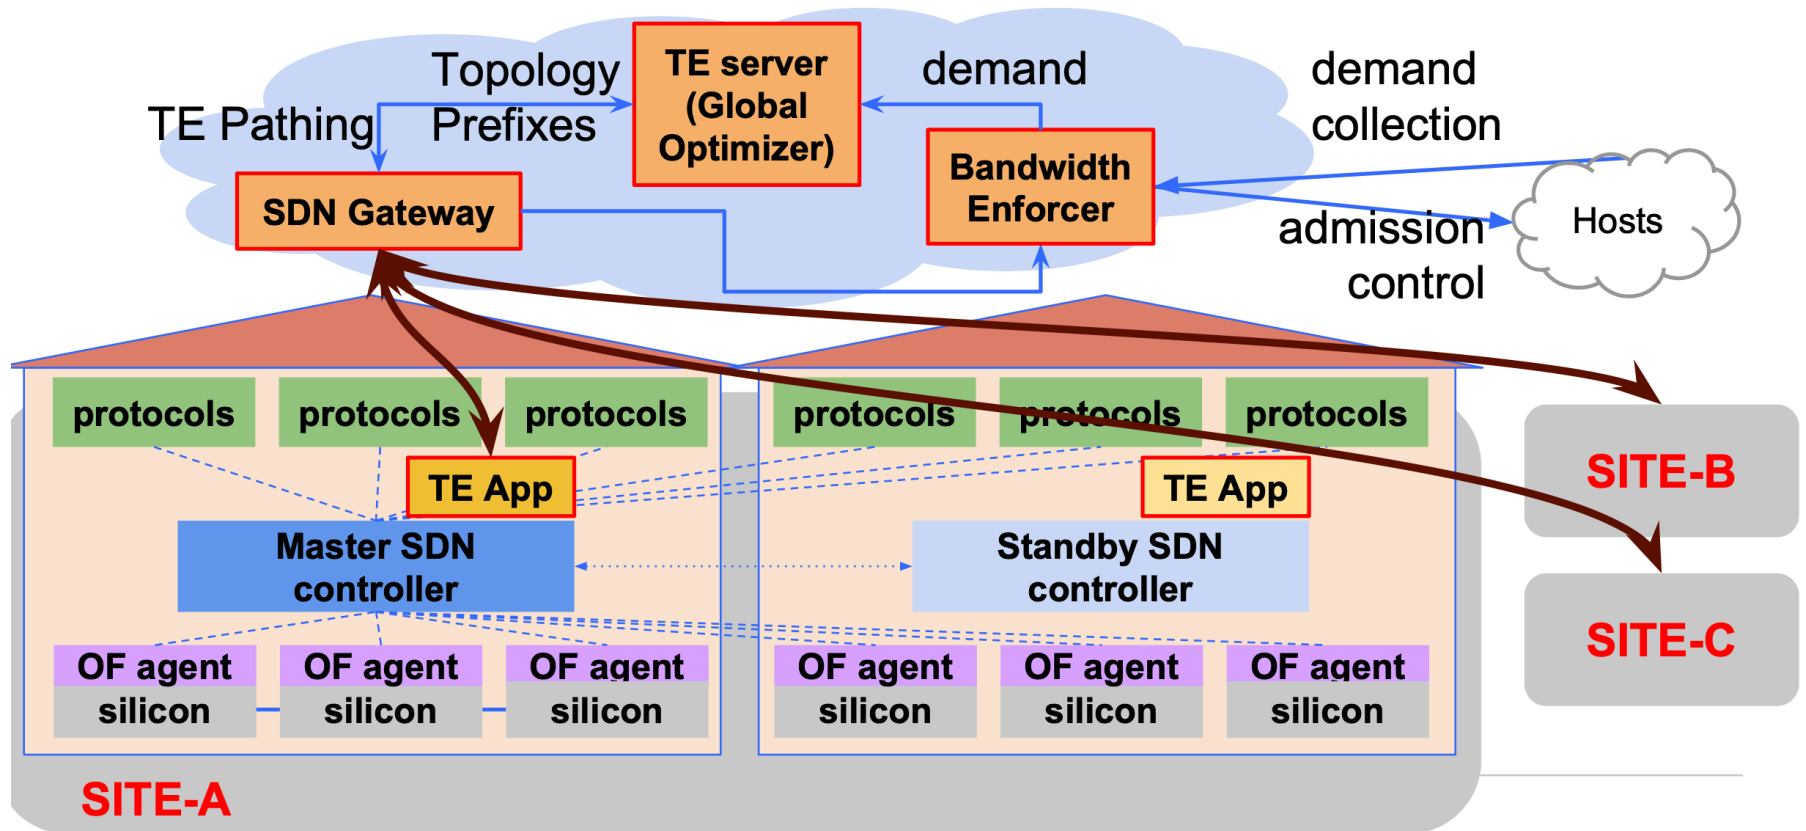


B4 SDN architecture




Unit of management is a site = fabric

B4 SDN architecture




Benefit of Centralized TE



- ~20% increase in throughput over SPF
- Larger benefits during capacity crunch

Lowers the requirement for bandwidth provisioning

Benefit of Centralized TE



B4 – your opinions

- Understandability of the paper:
 - Routing details were difficult to follow.
 - Quagga: routing protocol implementation on Linux.
 - TE algorithm was difficult to understand.
 - Objective: max-min fairness
 - A: 10Gbps, B: 5Gbps, total link capacity = 12Gbps
 - B = 5Gbps
 - A = 7 Gbps
 - A: 10Gbps, B: 5Gbps, C: 2Gbps, link capacity = 12Gbps
 - C = 2Gbps
 - B = 5Gbps
 - A = 5Gbps
 - Same demands, $W(A) = 2, W(B) = 1, W(C) = 1$, link capacity = 12Gbps
 - C = 2Gbps
 - B = 3.33Gbps
 - A = 6.67Gbps
 - *Bandwidth Enforcer, SIGCOMM'15 has more details on TE algorithms*

B4 – your opinions

- Pros:
 - Good example of use of OpenFlow
 - Nothing new and fancy, straight-forward application of OpenFlow.
 - Large-scale deployment, beyond campus networks
 - Concrete design
 - Cost budget
 - Considers single-point of failure / has a fault-tolerance mechanism
 - Aggregated TE – more scalable!
 - Able to achieve very high utilizations.
 - Real-deployment experiences (e.g. outage)

B4 – your opinions


- Cons:
 - Applicability to other WANs? Too specific to Google?
 - Does not work with commodity switches / needs custom hardware.
 - Net neutrality??
 - Why the greedy heuristic for TE? How close to optimal is it?
 - Why only 4 path choices?
 - “Why’s” not explained very well.
 - More details on failure handling needed.
 - What happens when an entire site goes down?
 - State consistency across control protocols not explained well.
 - Evaluation results over multiple days.
 - More example applications.

B4 – your opinions

- Ideas:
 - Minimize communication overhead between control and data plane.
 - More logging and monitoring, more route attributes (loss rates, delay, etc)
 - Analysis of TE solutions.
 - Better network availability guarantees.
 - Increased scalability.
 - Can ISPs provide more customized services to their customers?
 - What about Google's other WAN?


B4 and After: SIGCOMM'18

- Growth in traffic: more sites, larger sites, more paths.
 - Flat topology scales poorly:
 - Hierarchical topology at each site.
 - Hierarchical traffic engineering.



Another software-defined WAN


- SWAN (WAN connecting Microsoft's datacenter)
 - Goal: increase WAN link utilization.
 - Centralized and global traffic engineering.



Other SDN usecases at Google


Datacenter routing

- Few 100-1000 switches distributed across clusters.
- High communication overhead for distributed routing.
- Symmetric topology: multipath equal cost forwarding.




Datacenter routing

- Jupiter (Google's Datacenter)
 - Centralized configuration for baseline static topology.
 - Centralized dissemination of link state.
 - Each switch reacts locally to changes.



Policy enforcement at user-facing edge

- Internet edge routers implement rich set of features:
 - Access control, firewall, BGP routing policies.
- Policies require global, cross-layer optimizations.
 - Might also require switch upgrades, that affect availability.



Policy enforcement at user-facing edge

- Espresso:
 - Global software control plane to compute policies.
 - Local control plane to translate policy to forwarding rules.

