Bandwidth Isolation Guarantee for SDN Virtual Networks

Gyeongsik Yang · Yeonho Yoo · Minkoo Kang
Heesang Jin · Chuck Yoo
ksyang@os.korea.ac.kr

KOREA UNIVERSITY

IEEE INFOCOM
2021 IEEE International Conference on Computer Communications

SDN-based network virtualization (SDN-NV)

- Network virtualization: isolated network between VM and containers (entities) of tenants
- SDN-NV allows programmable virtual networks in datacenters

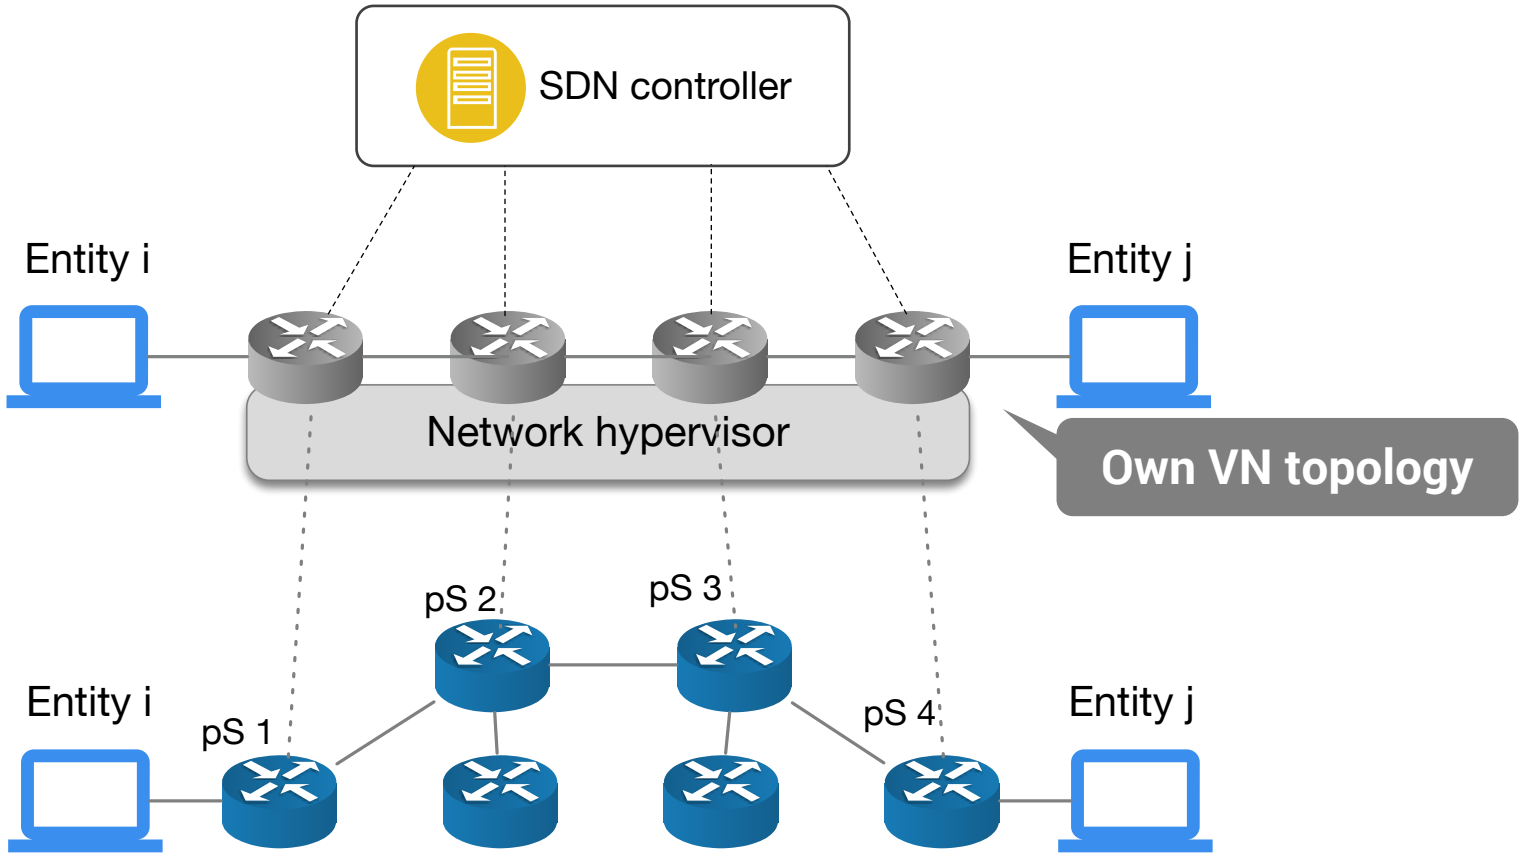

SDN-based network virtualization (SDN-NV)

- Custom packet forwarding path through SDN controller of tenants
- Path installed by individual flow rules

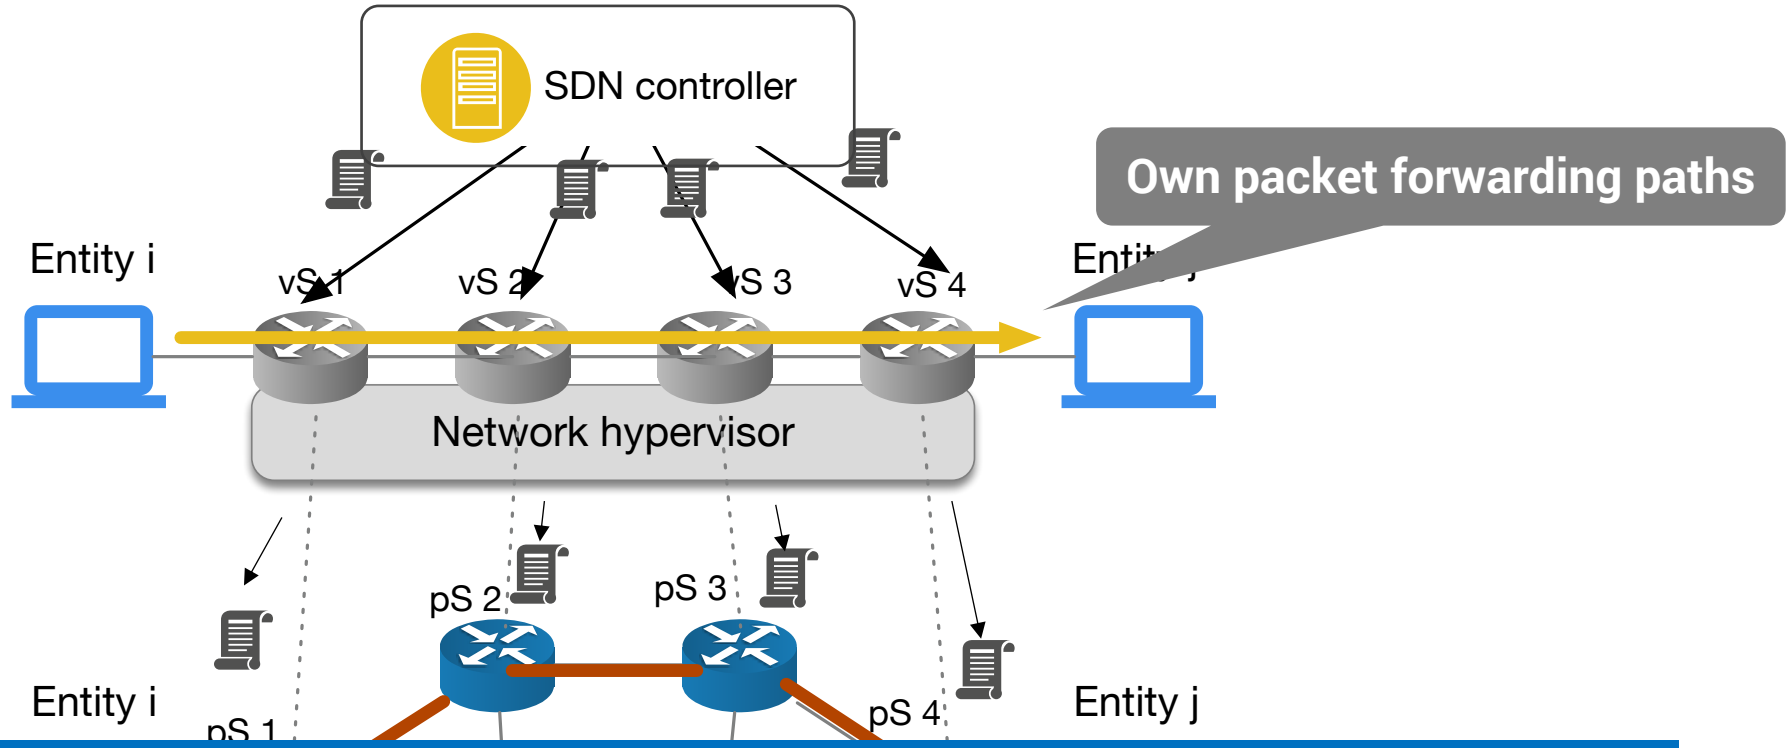

Problem on Bandwidth Isolation Guarantee

Bandwidth isolation guarantee

- Bandwidth requirement for each datacenter application
  - Significant revenue loss and service quality downgrade when the requirement is not met
- Bandwidth isolation guarantee for
  1) keeping the demanded bandwidth; 2) independent to other tenants

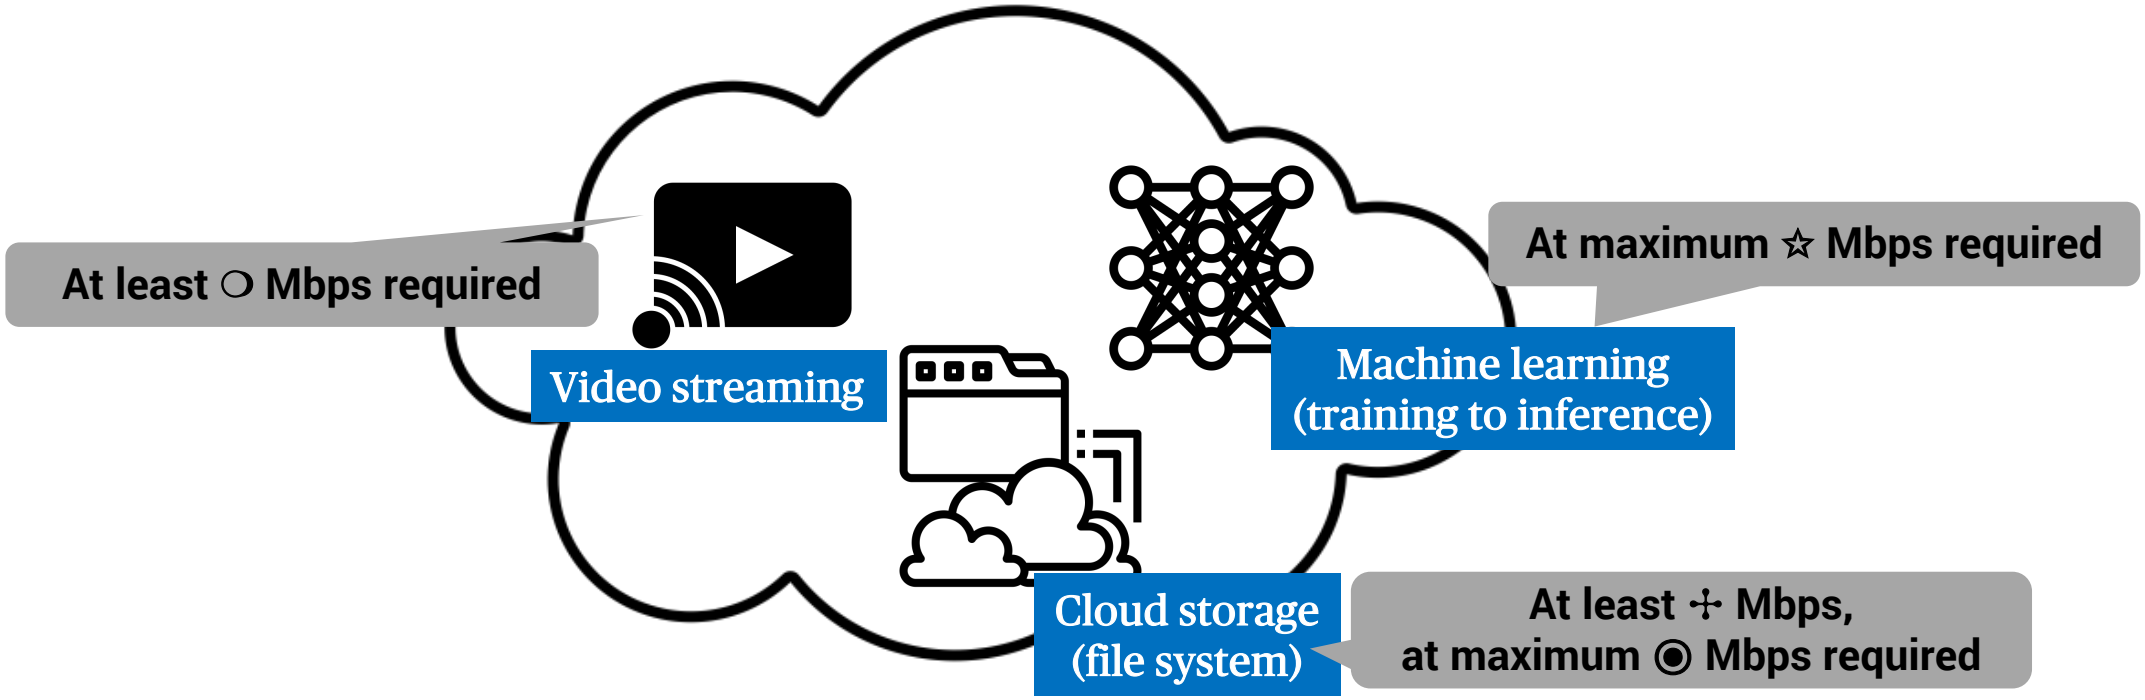

Lack of bandwidth isolation guarantee in SDN-NV

	FlowVisor (OSDI 2010)	FlowN (IEEE Internet Comput. 2012)	OpenVirteX (HotCloud 2014)	CoVisor (NSDI 2015)	OnVisor (IJNM 2018)	Libera (IEEE Commag 2020)	V-Sight (IEEE INFOCOM 2020)
Multi-tenancy	O	O	O	X	O	O	O
Architecture	Central	Central	Central	Central	Distributed	Central	Central
Scalability	X	X	X	X	O	O	O
VM migration support	X	X	X	X	X	O	X
Network monitoring	X	X	X	X	X	X	O
Bandwidth isolation guarantee	**X**	**X**	**X**	**X**	**X**	**X**	**X**

Bandwidth isolation guarantee studies on other NV technologies

- Mostly abstracted as a hose model
  - Bandwidth requirements given for non-blocking switch
  - Routing (e.g., cost-based shortest path) from datacenter operators

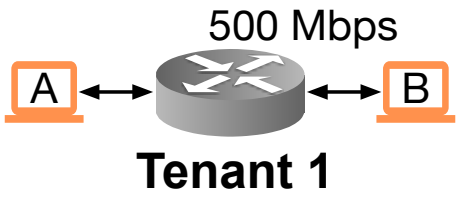

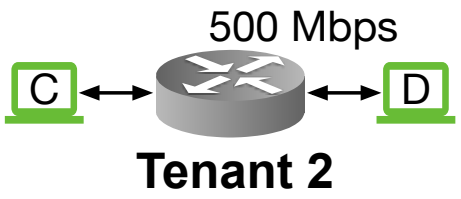

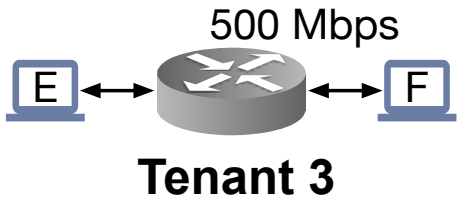

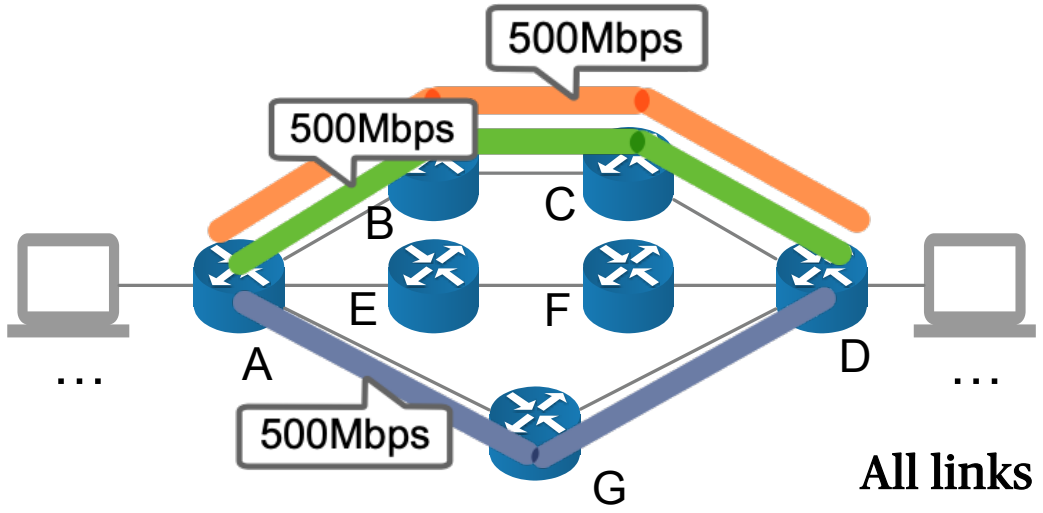

All links and switches: 1 Gbps

Detailed comparison and analysis on existing studies in paper

Problem of existing studies on SDN-NV

- Tenants perform routing; consequently, the links of paths can be overloaded

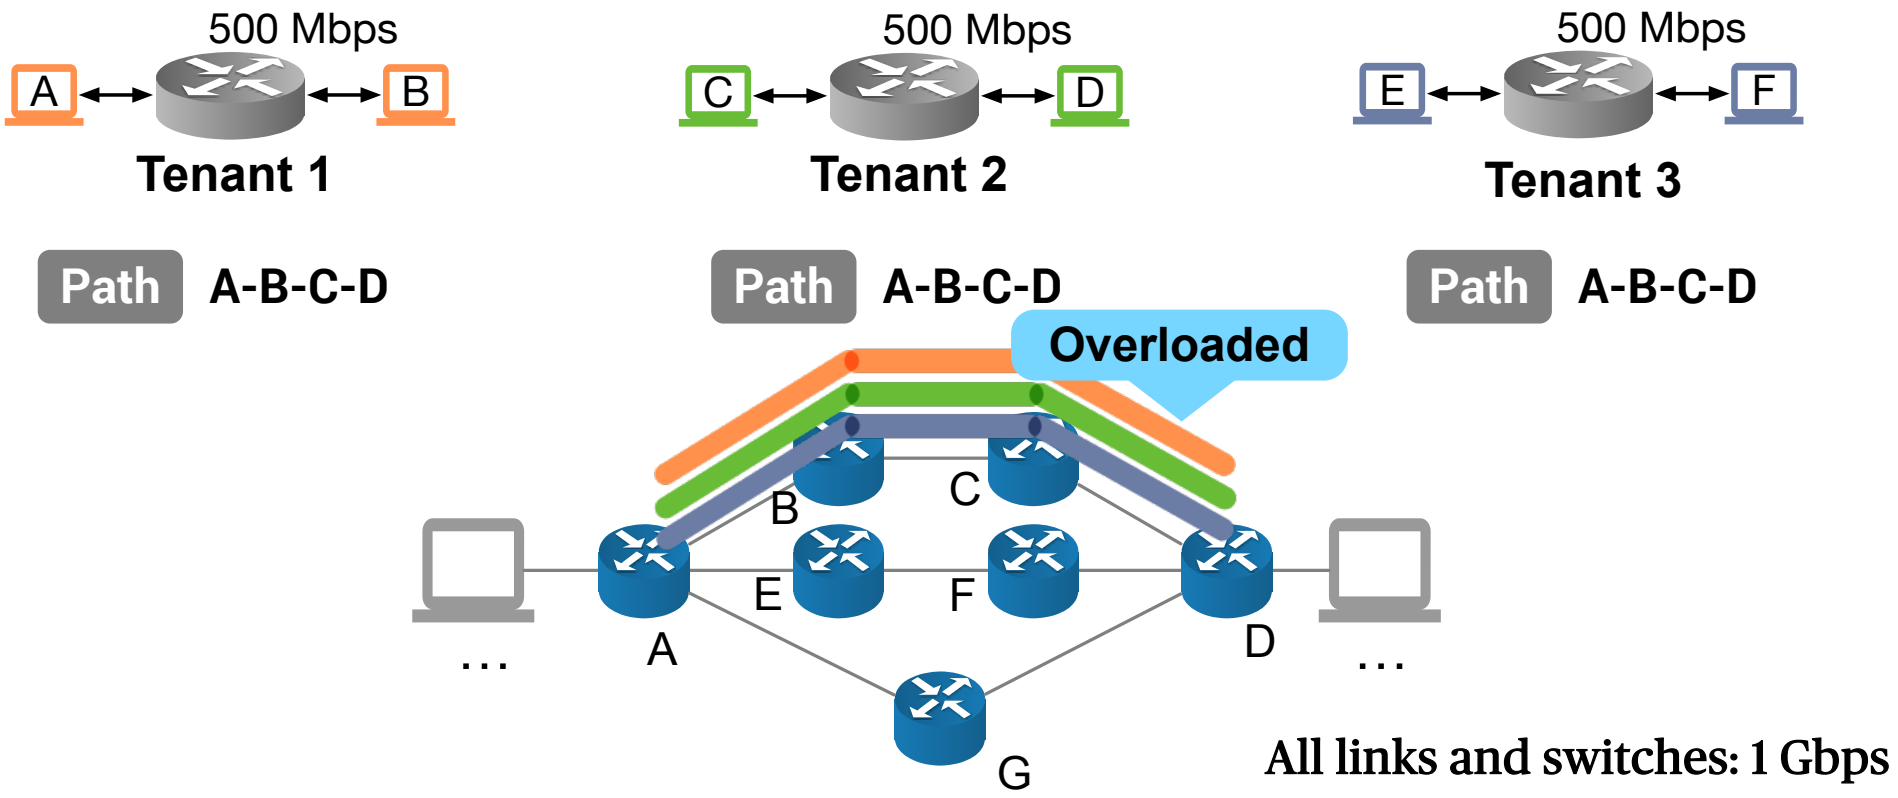

Solution: TeaVisor

- Preservation of paths from tenants by association with bandwidth requirements
- Multipath routing for bandwidth isolation guarantee

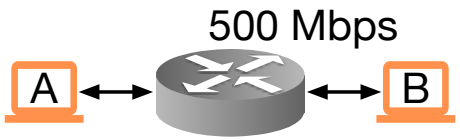

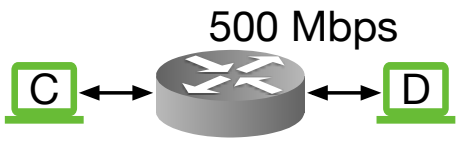

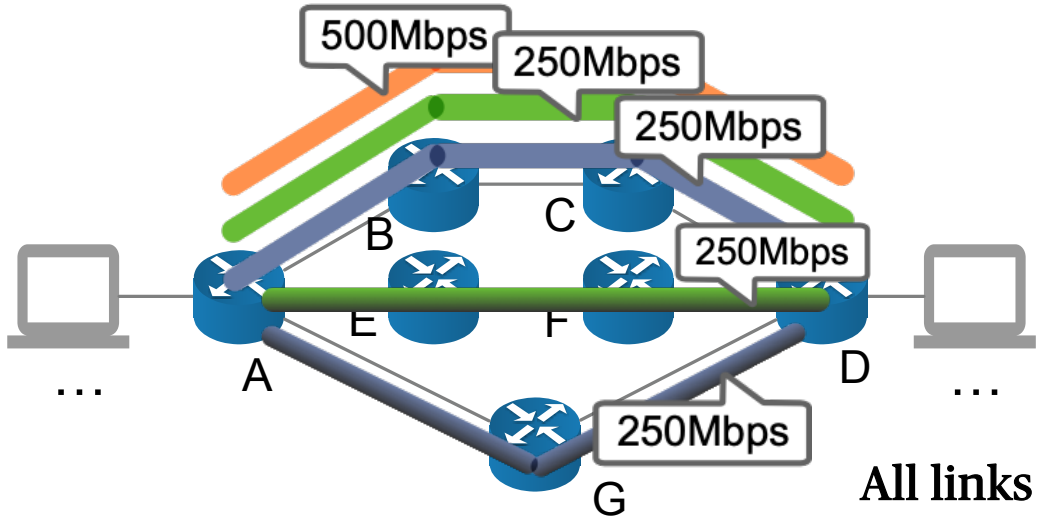

All links and switches: 1 Gbps

TeaVisor—three components

- **Path virtualization**
  - Multipath routing based on virtual path from tenants
- **Bandwidth reservation**
  - Per-path ”minimum” and ”maximum” bandwidth isolation guarantees
- **Path establishment**
  - Enforcement of multipaths and reserved bandwidths on the paths

Workflow (1) initiation

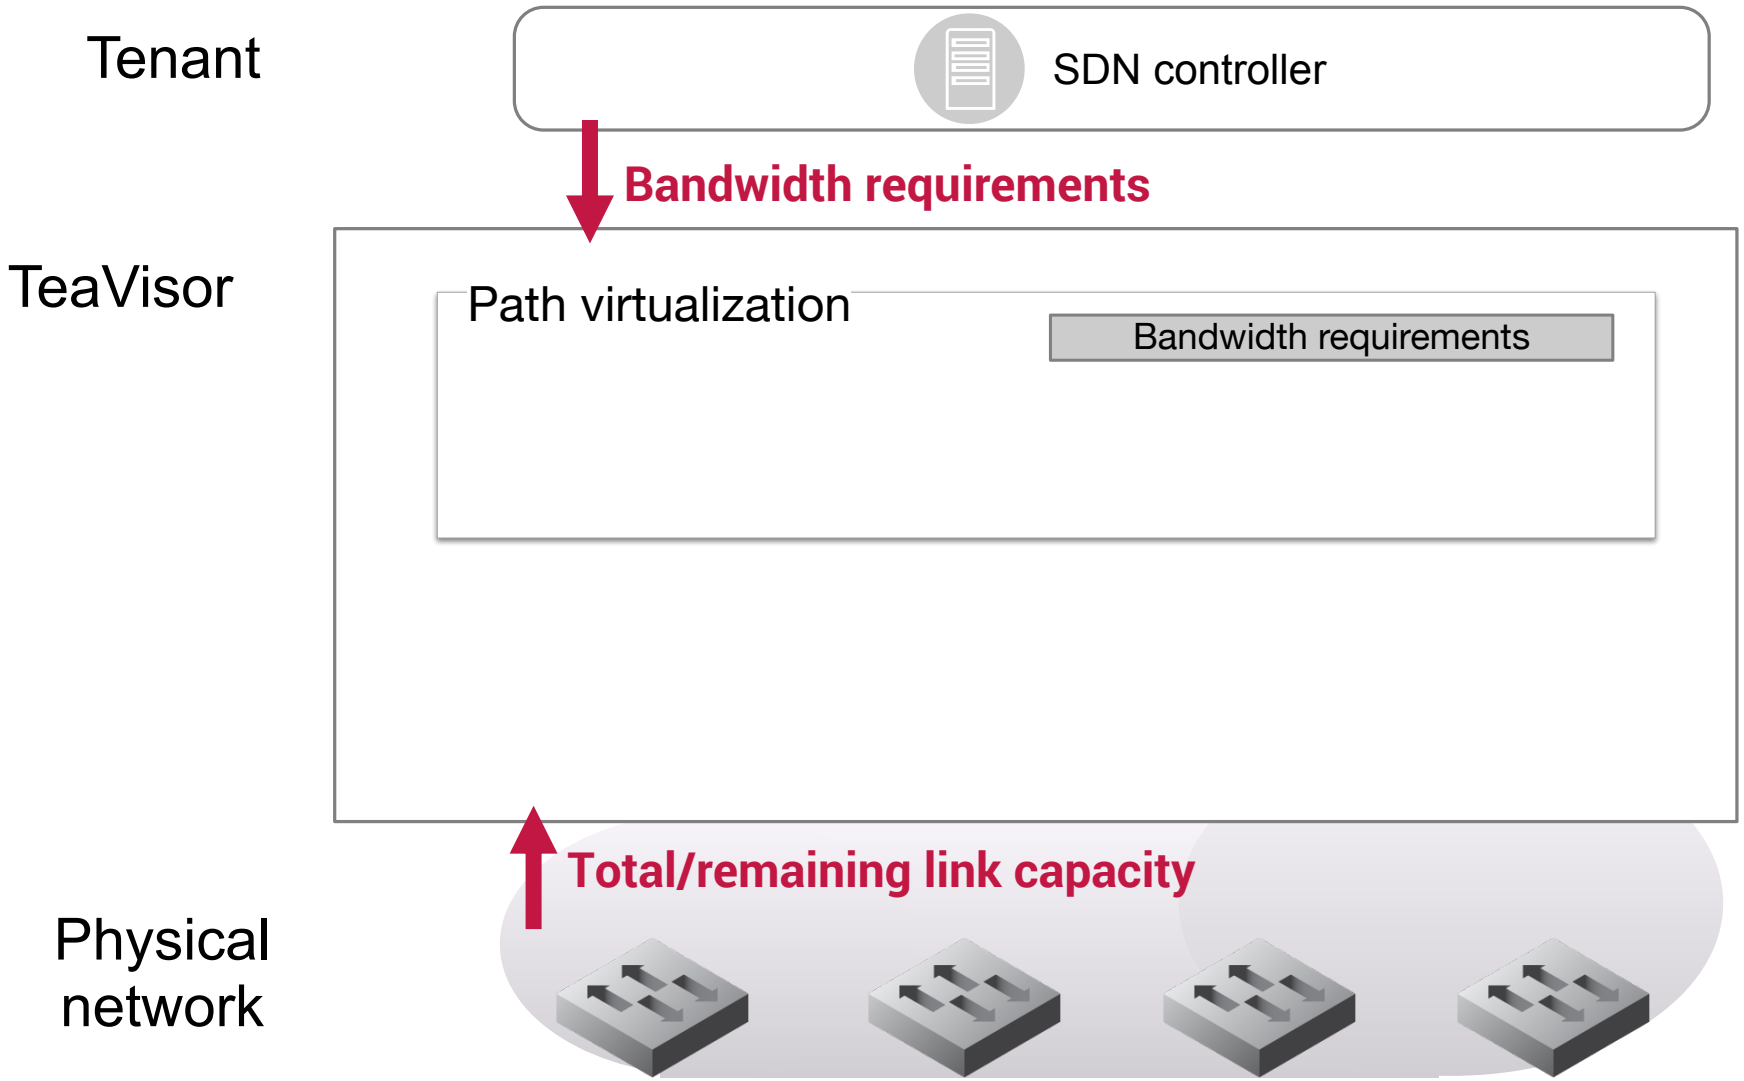

Workflow (2) flow rule generation from SDN controller

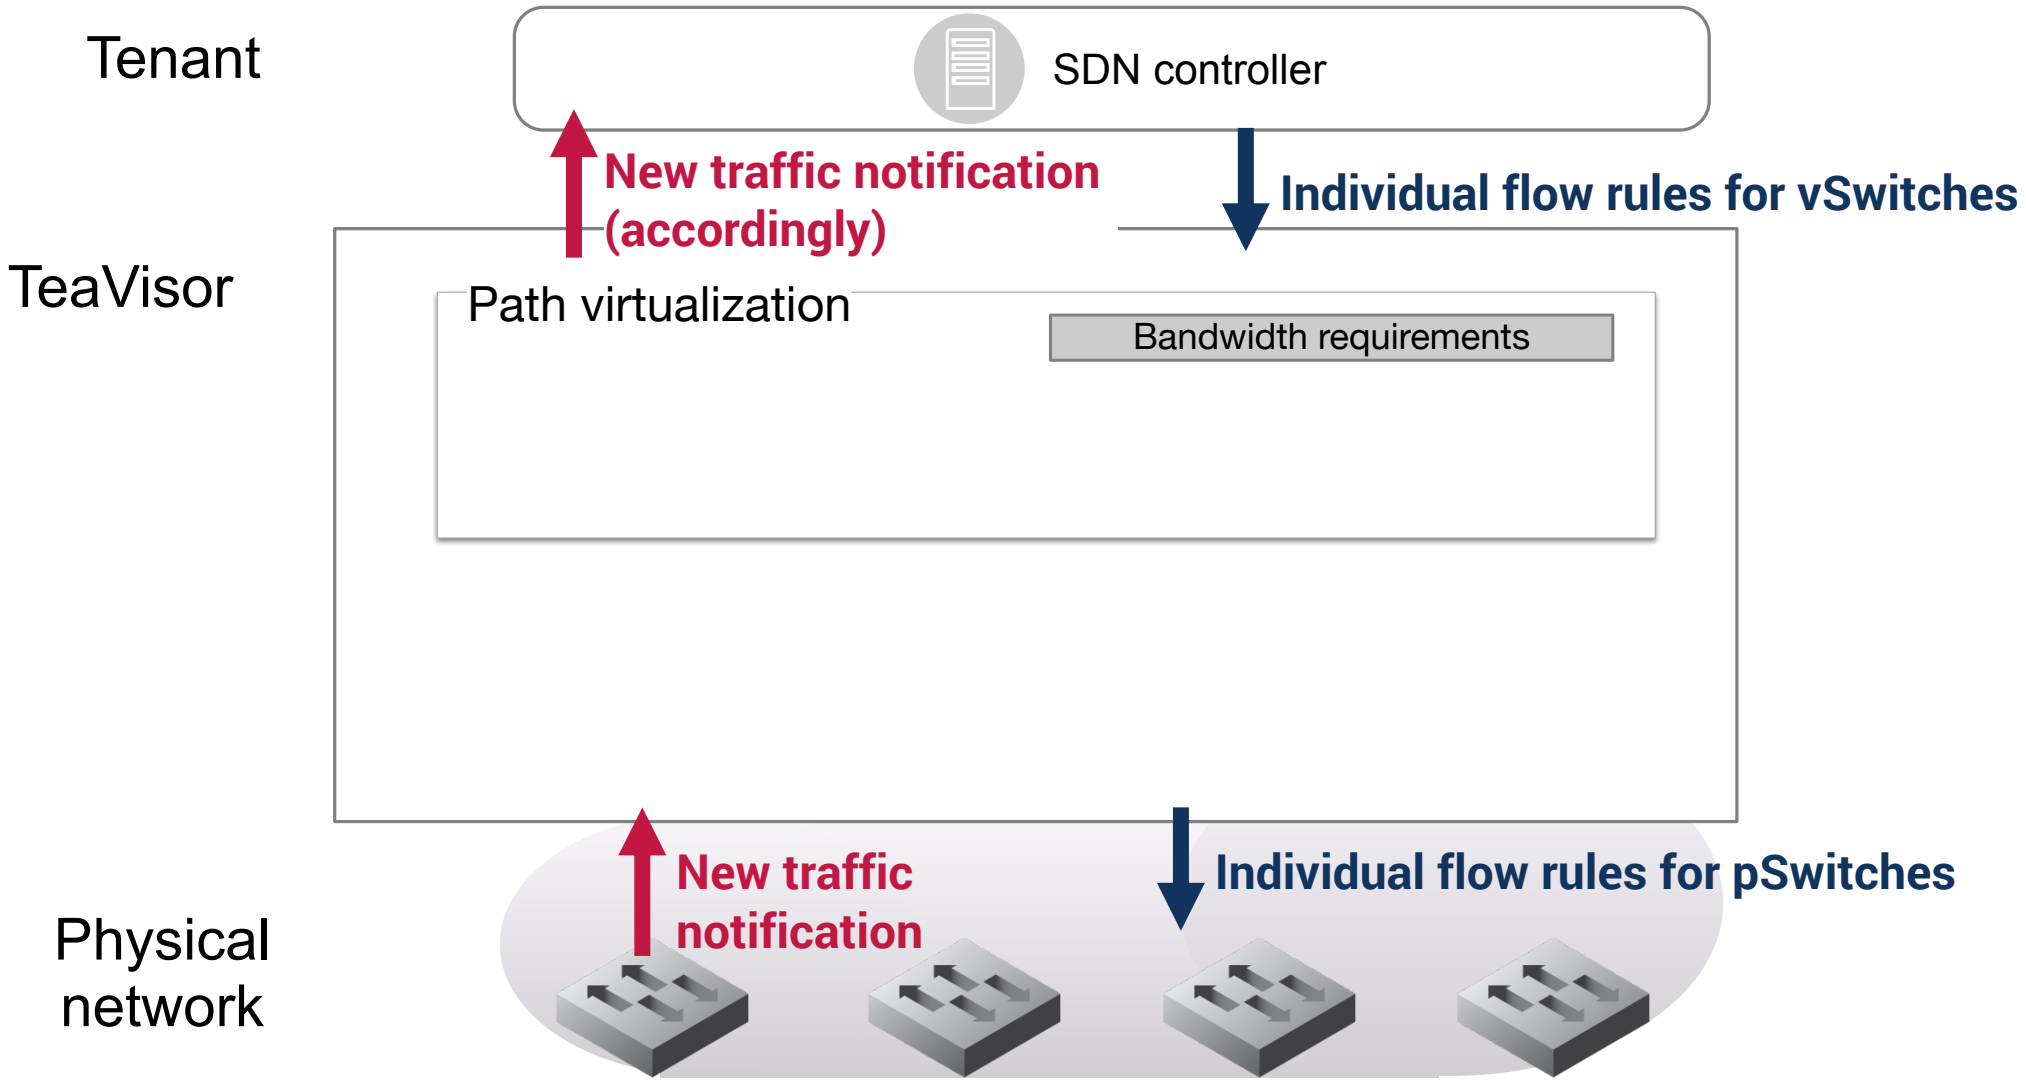

Workflow (3) path virtualization

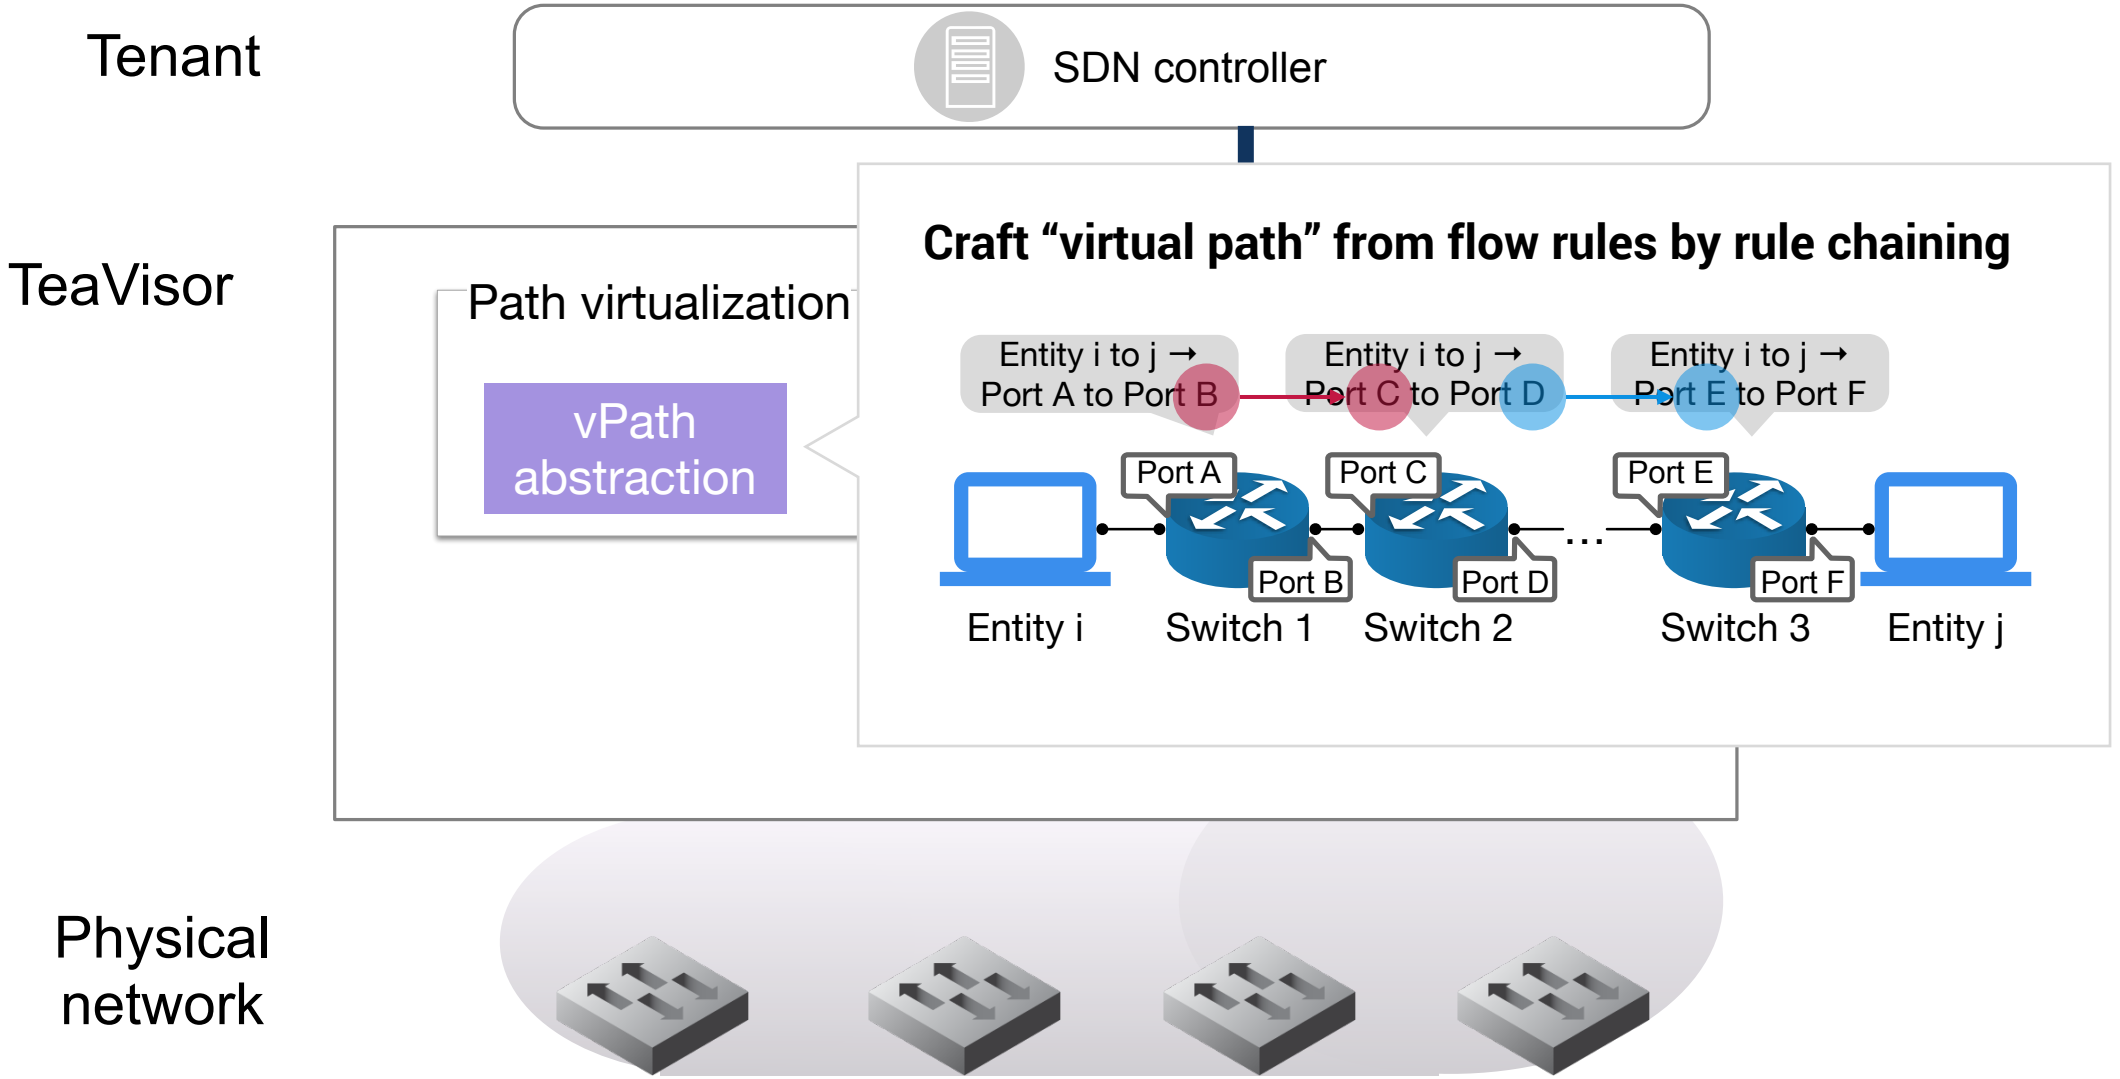

Workflow (3) path virtualization

Generation of physical paths

- Main path (mPath): generated by topology mapping
- Extra path (exPath): additional path for bandwidth requirement

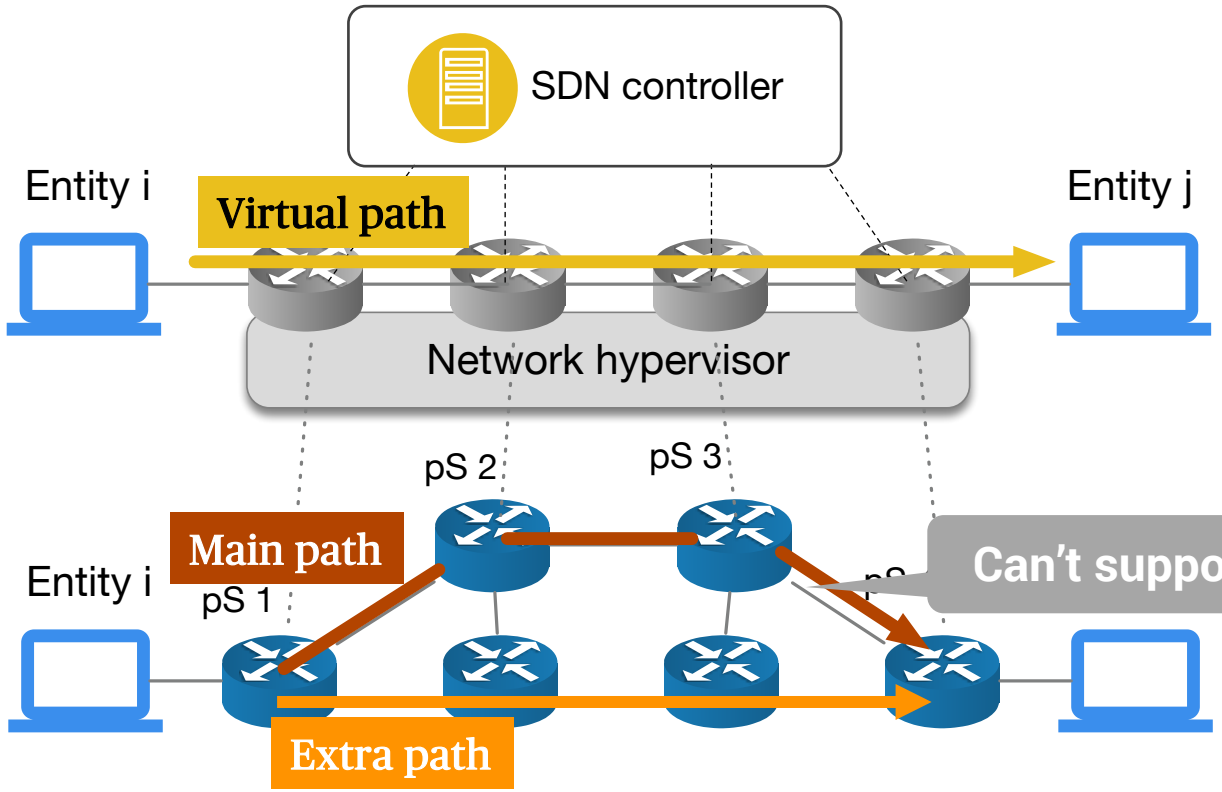

troller

Individual flow rules for vSwitches

width requirements

Multipath derivation

Can’t support bandwidth requirement alone?

Workflow (3) path virtualization

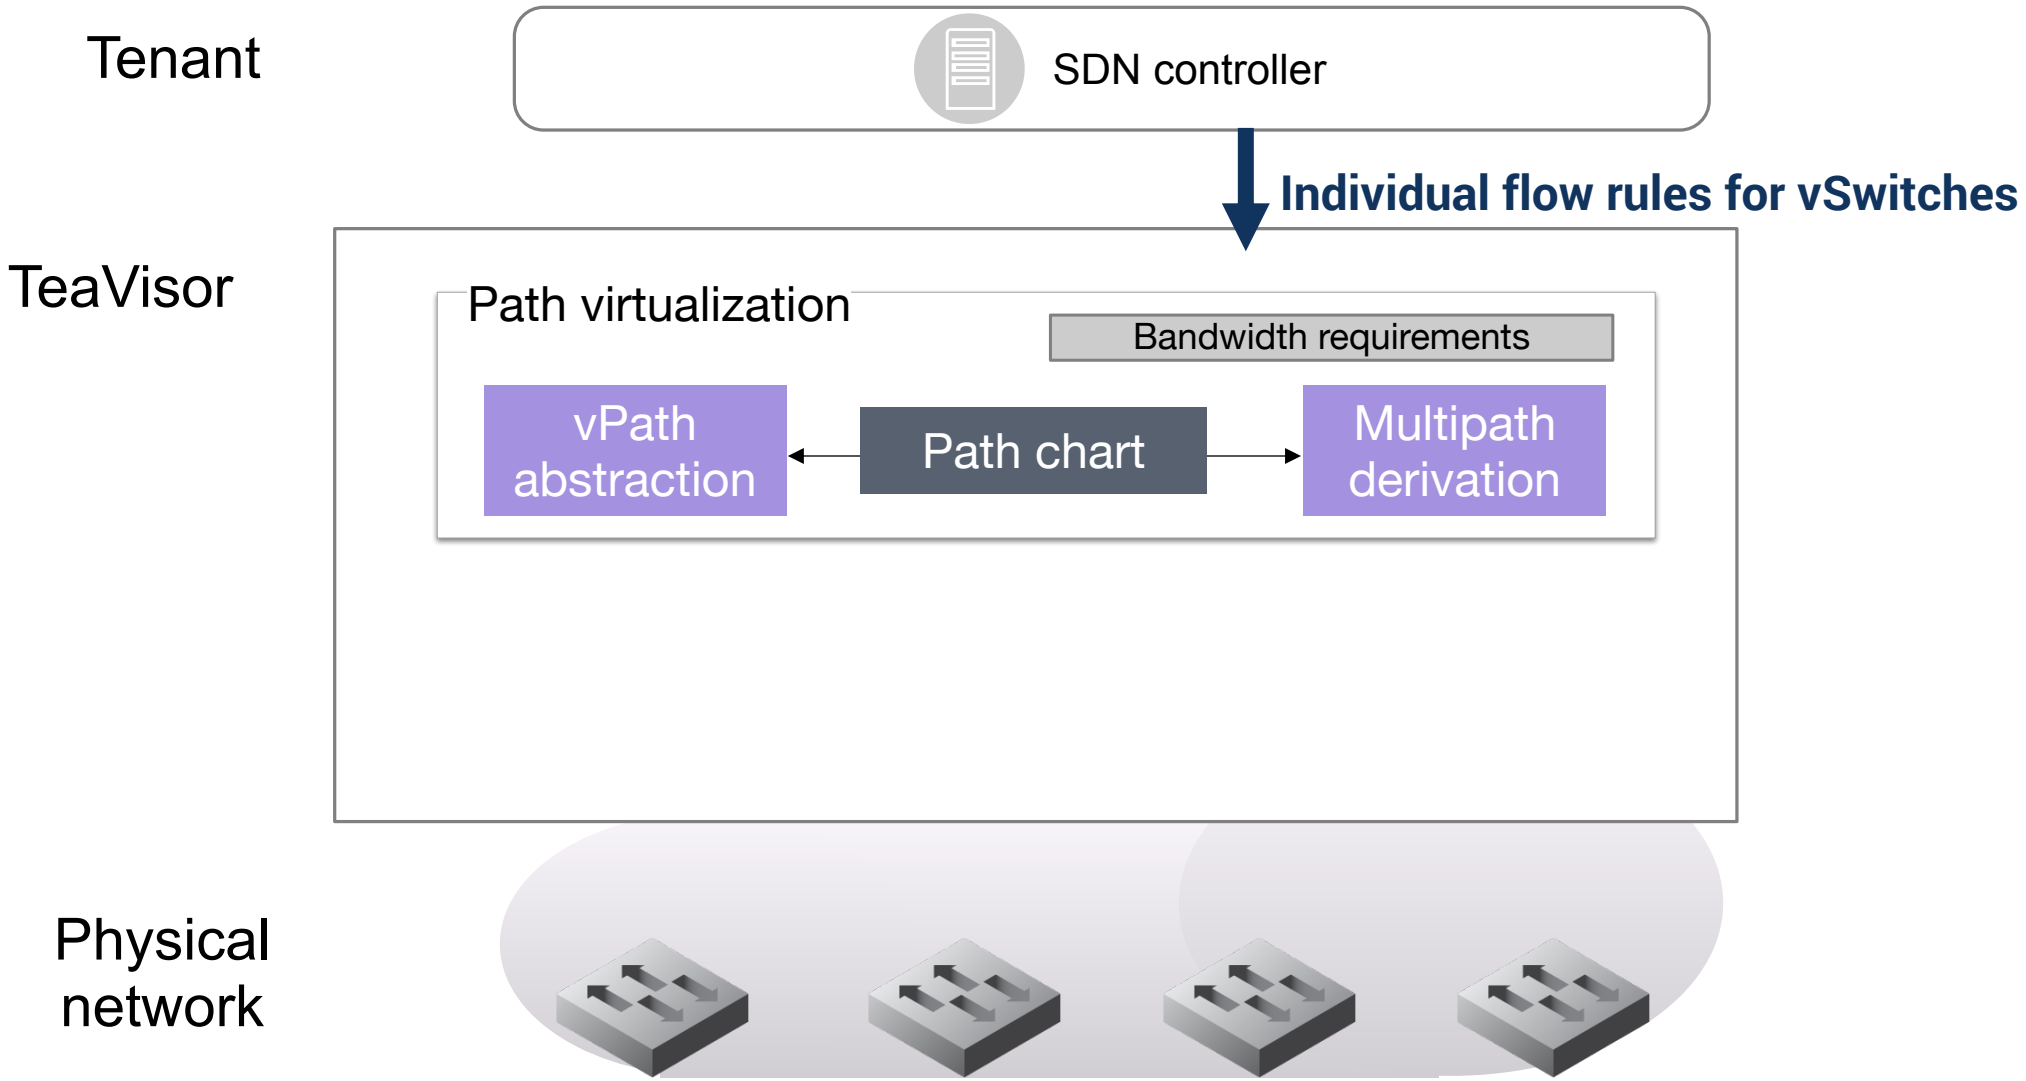

Workflow (4) bandwidth reservation

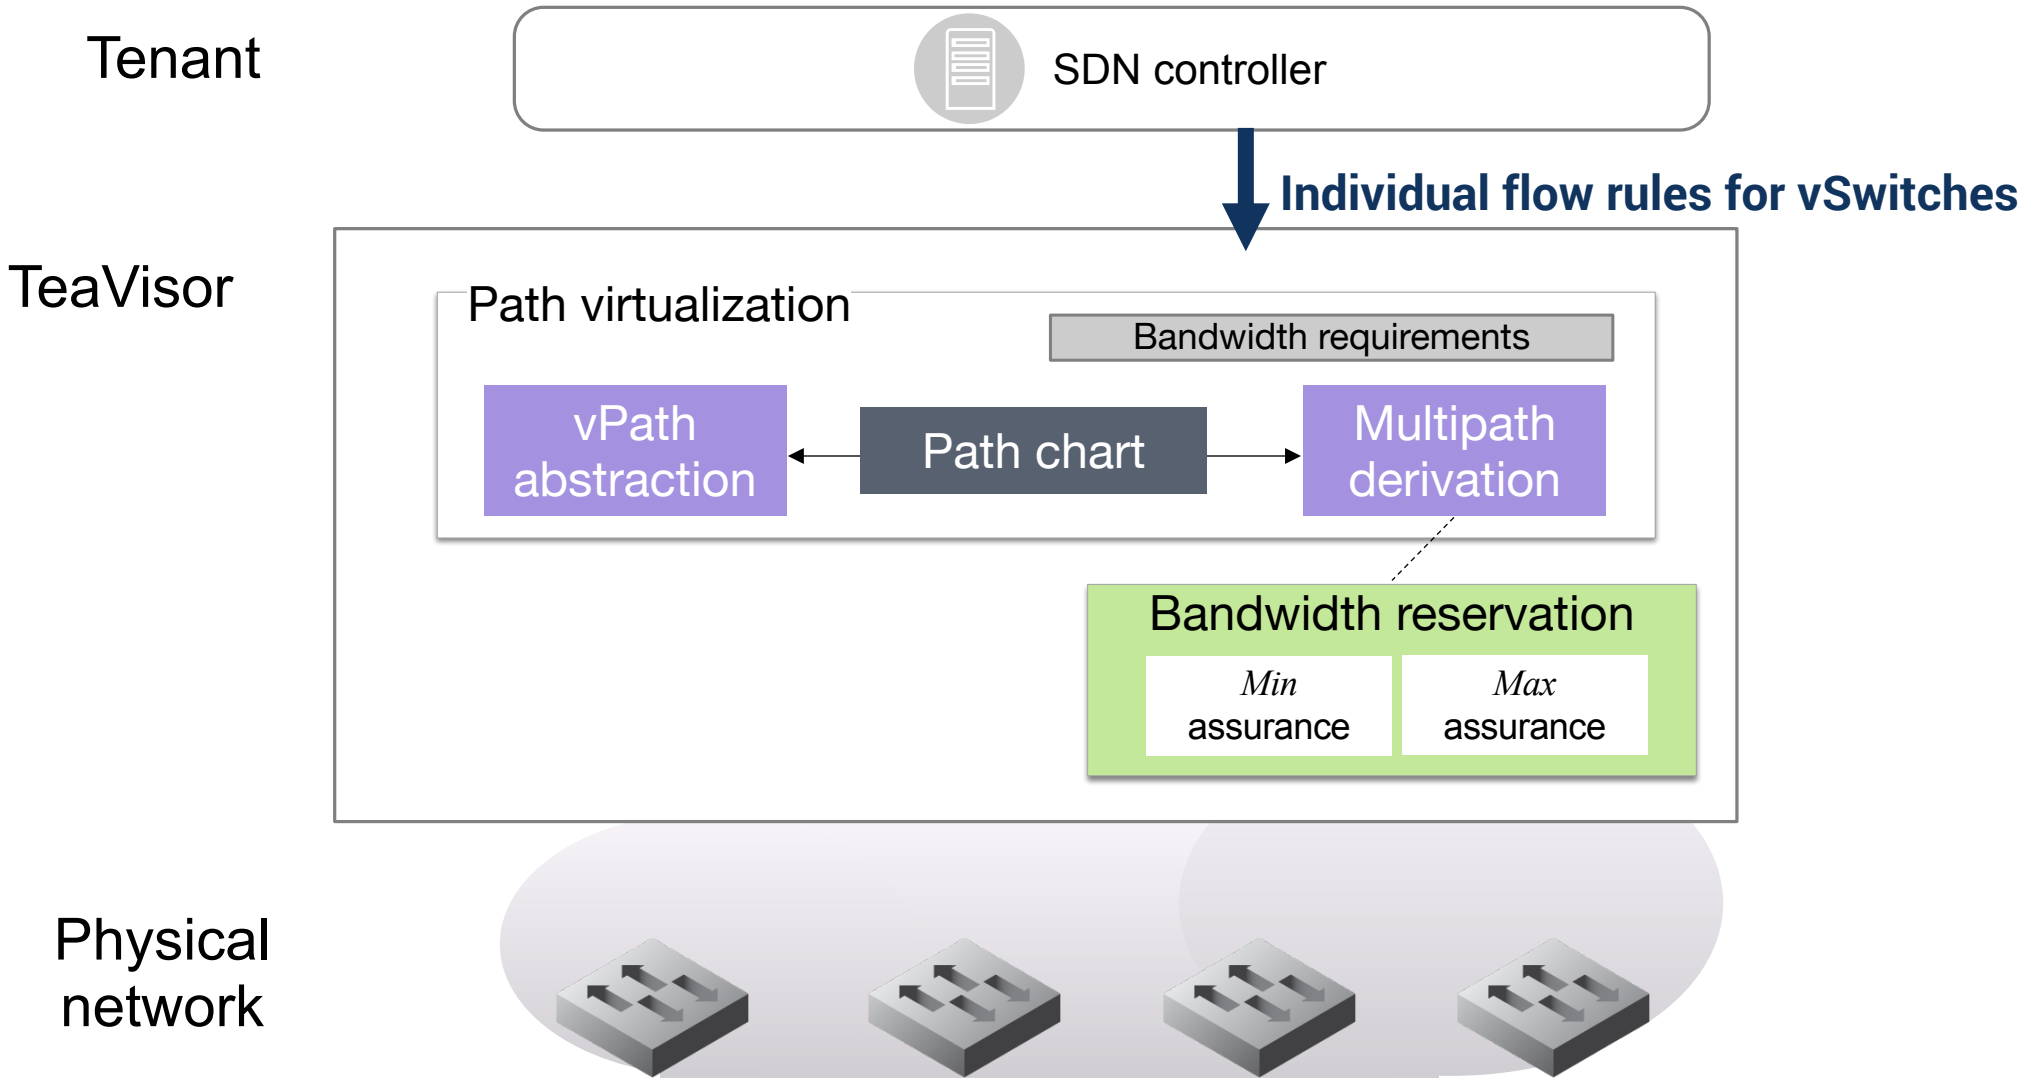

Bandwidth reservation – *Min* assurance

- Minimum requirement (*Min*): always reserved for a path

Entity pair	Pair1	Pair2
Min	800 Mbps	800 Mbps
Main path	A-B-C-D	A-G-D

All links and switches: 1 Gbps

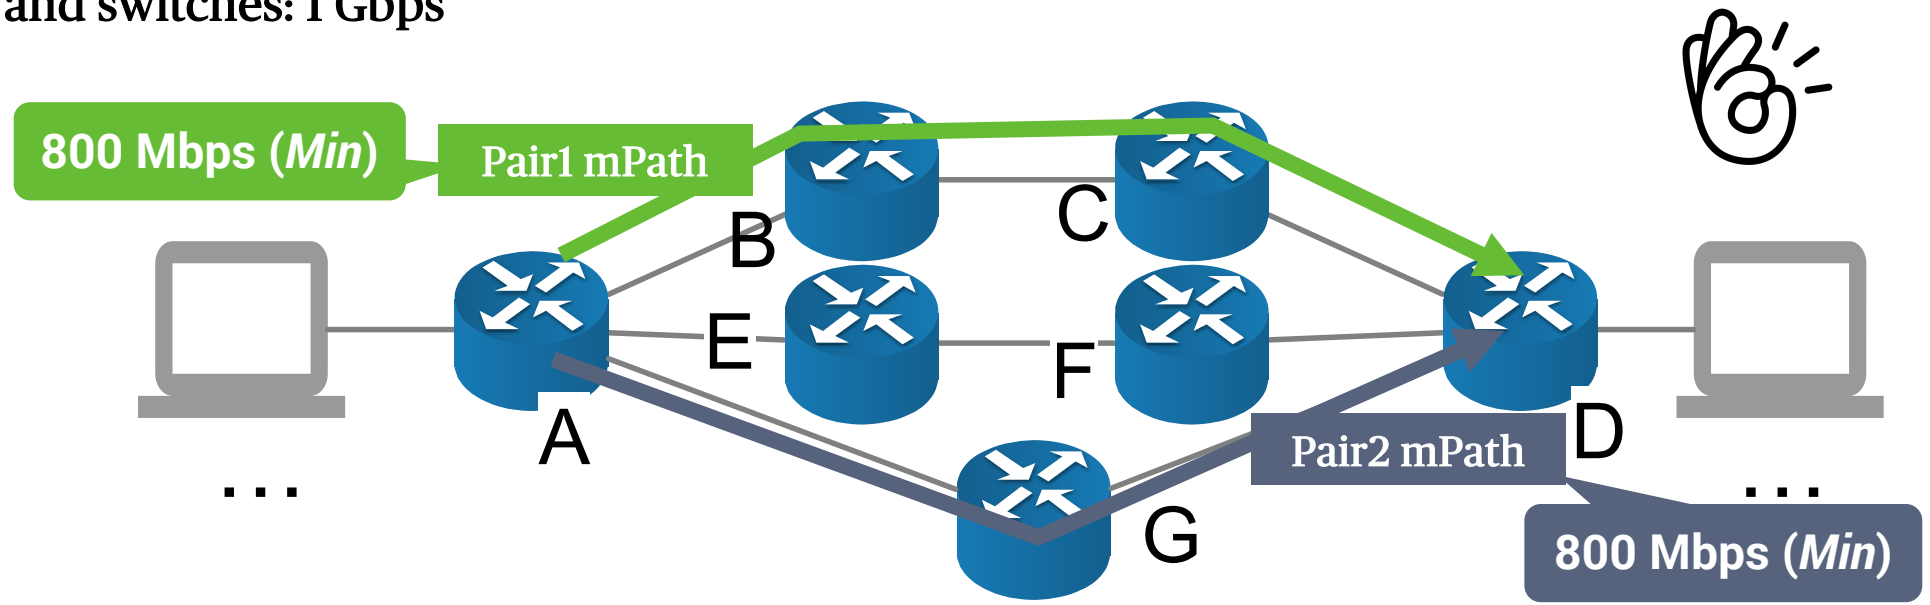

Bandwidth reservation – *Min* assurance

- The amount of bandwidth to be reserved on main path

Entity pair	Pair1	Pair2	Pair3
Min	800 Mbps	800 Mbps	800 Mbps
Main path	A-B-C-D	A-G-D	A-B-C-D

All links and switches: 1 Gbps

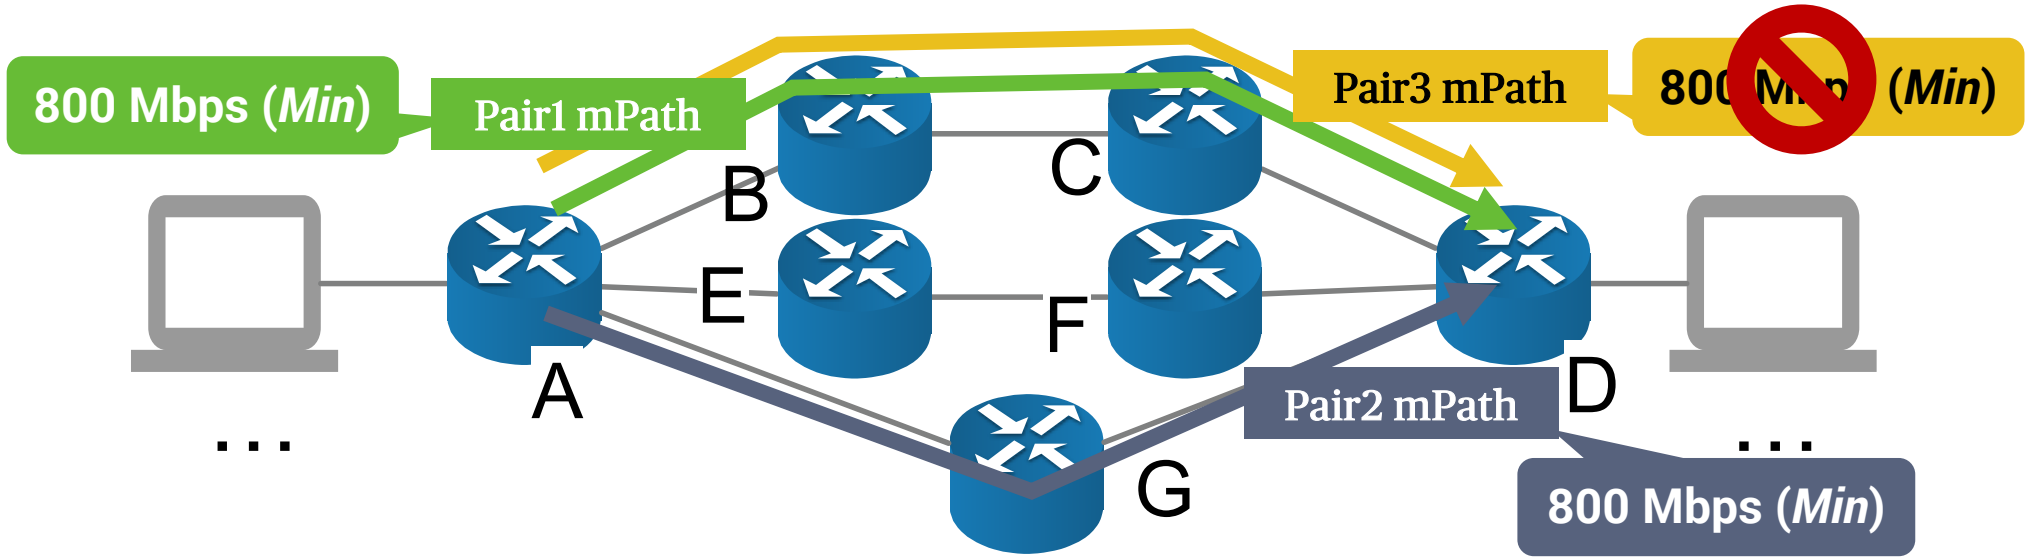

Bandwidth reservation – *Min* assurance

- Reserve *Min*d* on mPath (*d*: bandwidth reservation ratio)

Entity pair	Pair1	Pair2	Pair3
Min	800 Mbps	800 Mbps	800 Mbps
Main path	A-B-C-D	A-G-D	A-B-C-D

All links and switches: 1 Gbps

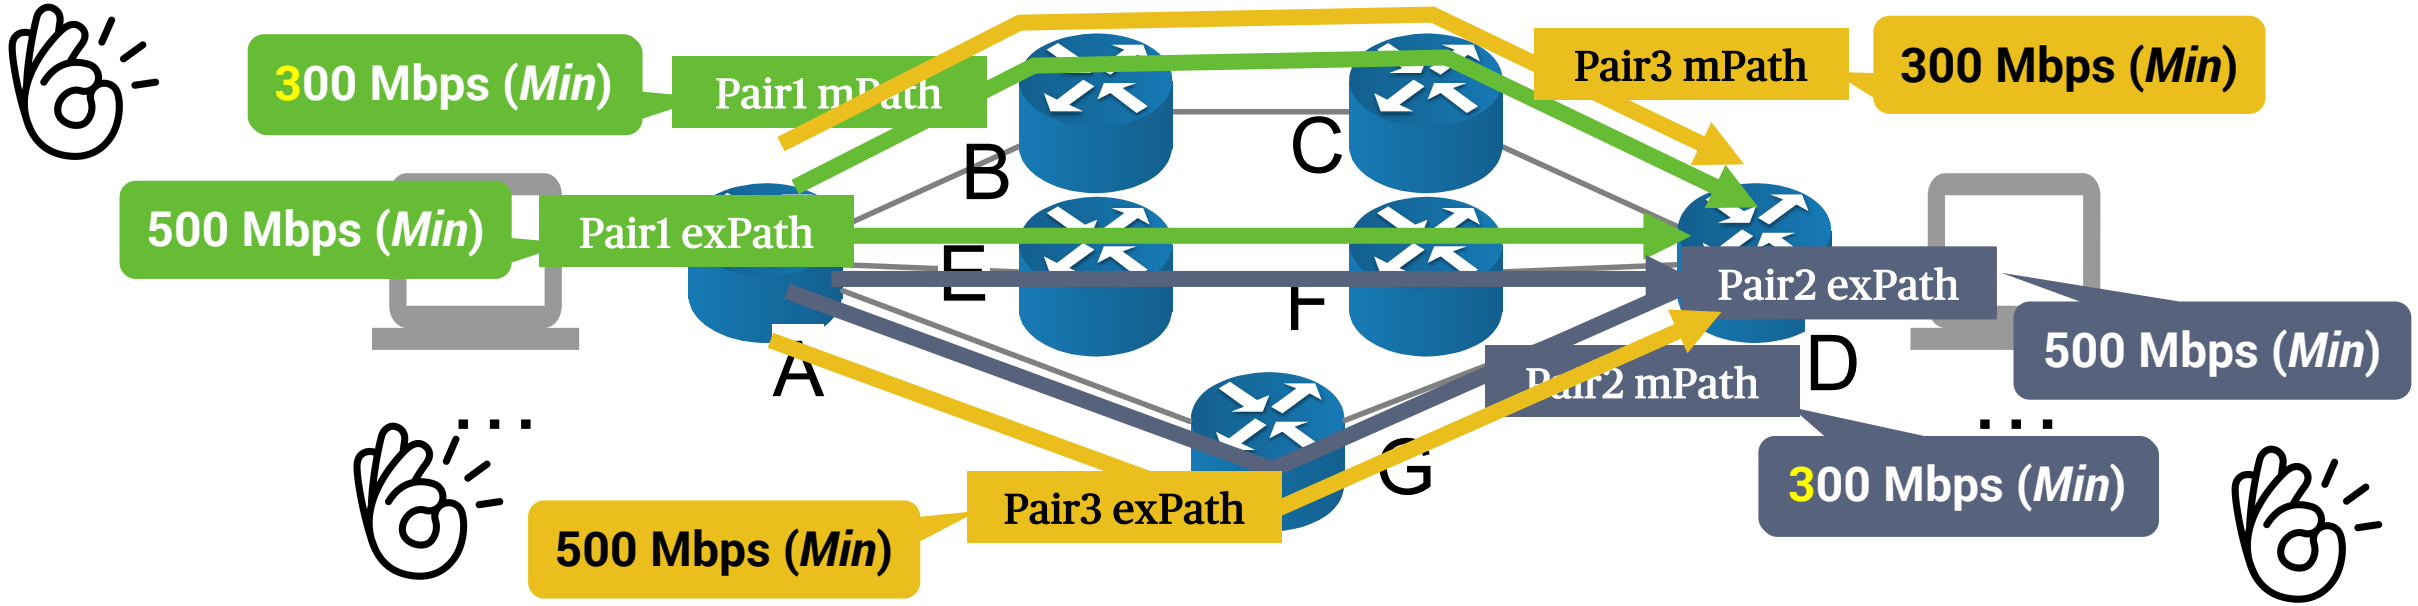

Min assurance - summary

- 1) Main path on *Min*d*
  - If the main path has capacity for Min*d → reserve for it
  - Otherwise, reduce bandwidth reserved for *Max* (next slide)
  - If not all possible, reject the pair on the mPath (access control)

- 2) Find extra path disjoint with main path
  - Reserve remaining *Min* repeatedly until the Min becomes zero

- 3) No available capacity on physical network → reject (access control)

- Detailed operation and policy in paper

Bandwidth reservation – *Max* assurance

- Increase of bandwidth for the existing pPaths at a regular interval
  - Detailed algorithms in paper
- Policies

Divide idle bandwidth within the *Max* bound

Equal share (ES)
equally between physical paths

Proportional share for bandwidth requirements (PBR)
according to the ratio of *Min* between physical paths

Proportional share for actual network usage (PNU)
based on actual throughput used within the reserved bandwidth (monitoring results)

Workflow (5) path establishment

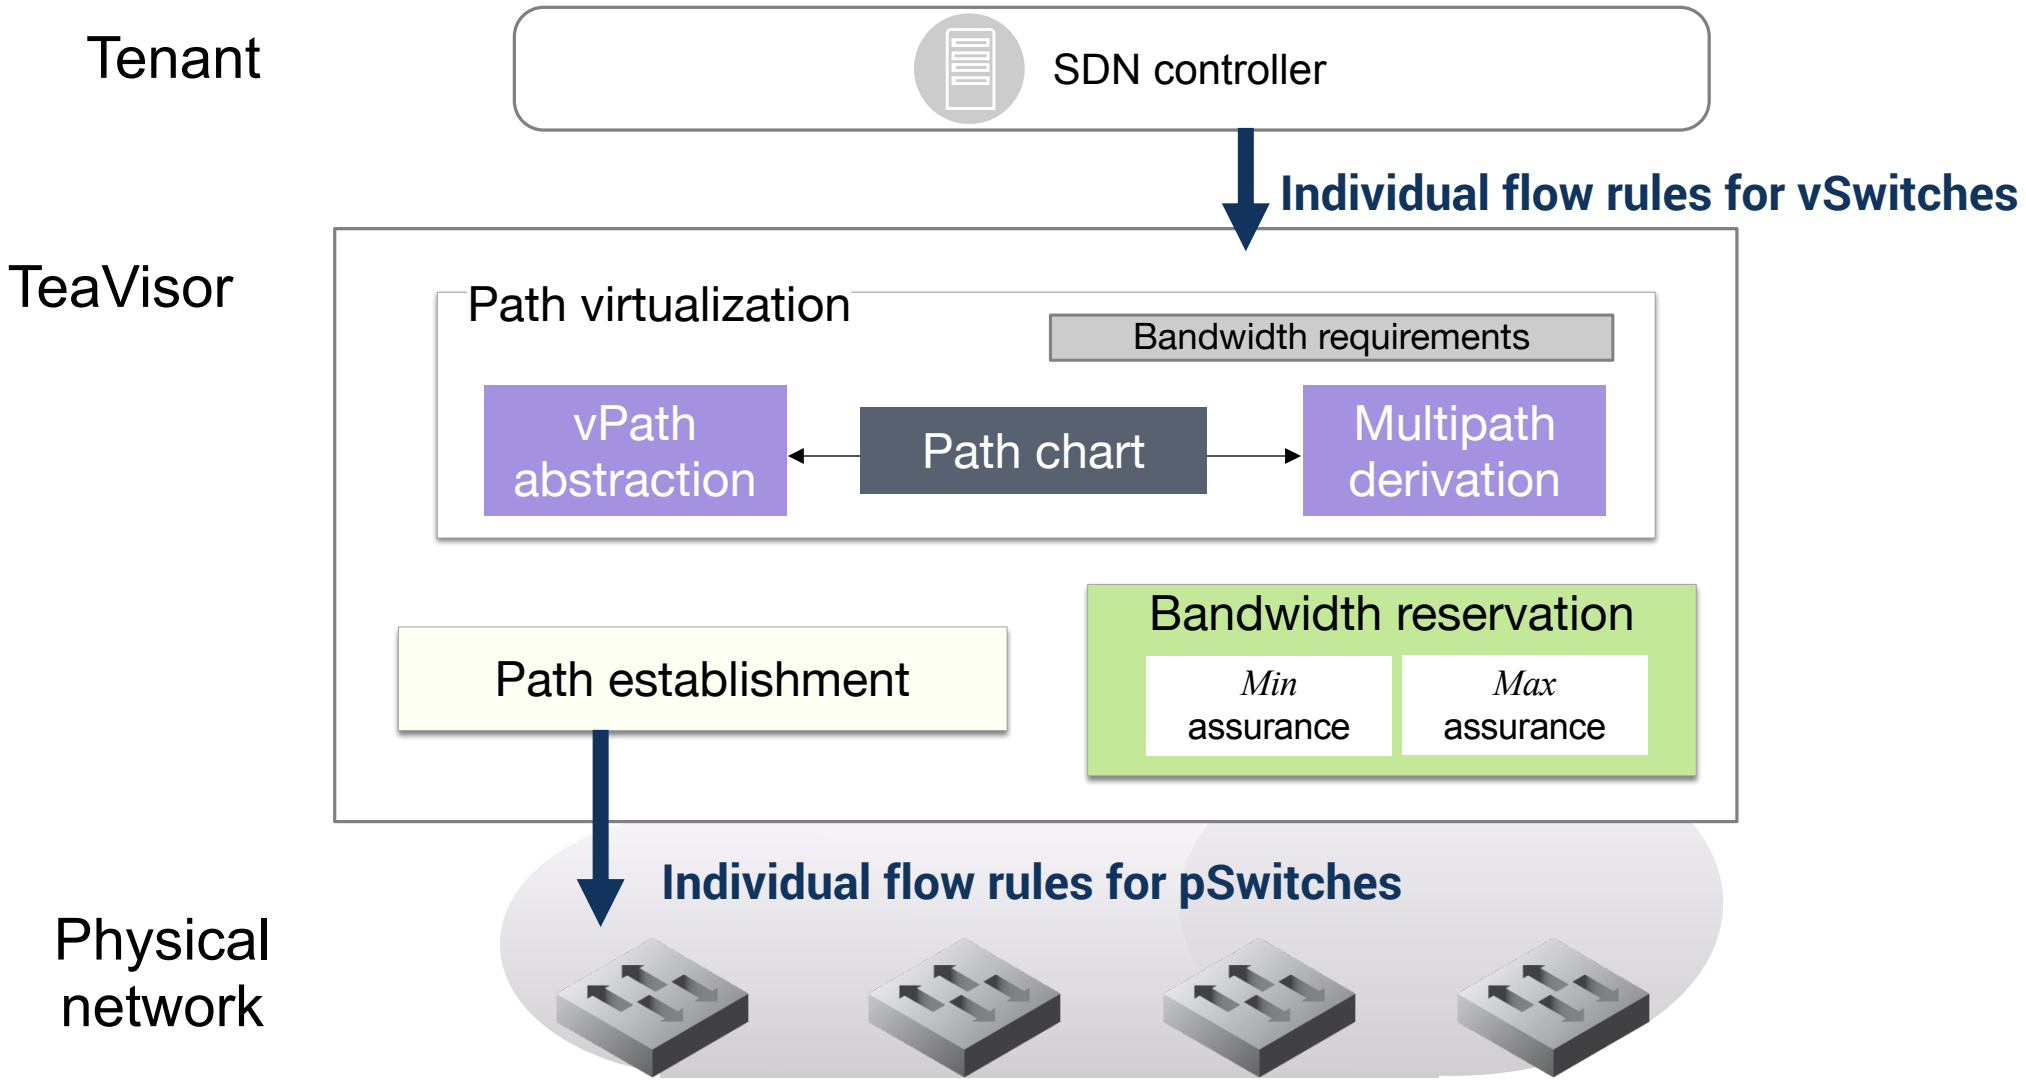

Evaluation settings

- **Comparisons**

  1) Libera: latest open-source SDN-NV solution

  2) Libera+RL (rate-limiting): SDN-NV + existing bandwidth isolation guarantee study ; rate-limiting on main path as its *Min* without multipath routing

  3) **TeaVisor**

- **Network Environment**
  - Topology: 4-ary fat-tree (Open vSwitch)
  - One to four tenants, four entity pairs per tenants
    - Each entity pair: 128 TCP connections

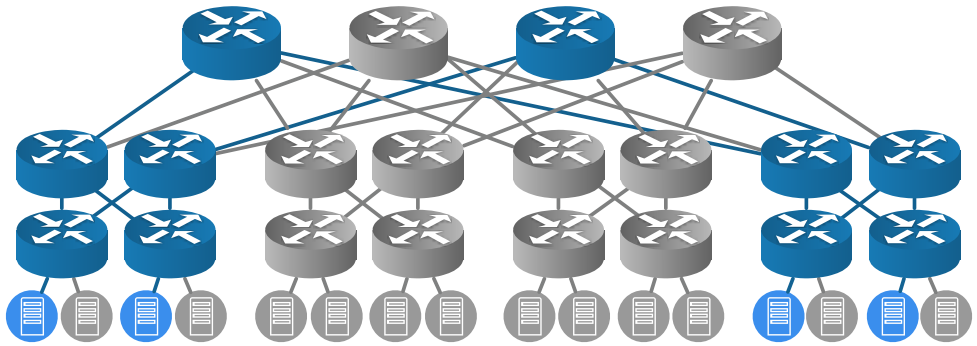

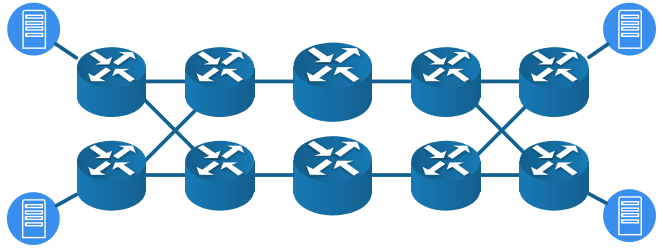

Bandwidth isolation guarantee

- Four tenants with 16 entity pairs
  - Sum of *Min* for all entity pairs: 75% of entire physical network capacity
  - Max of each entity pair: 125% of *Min*

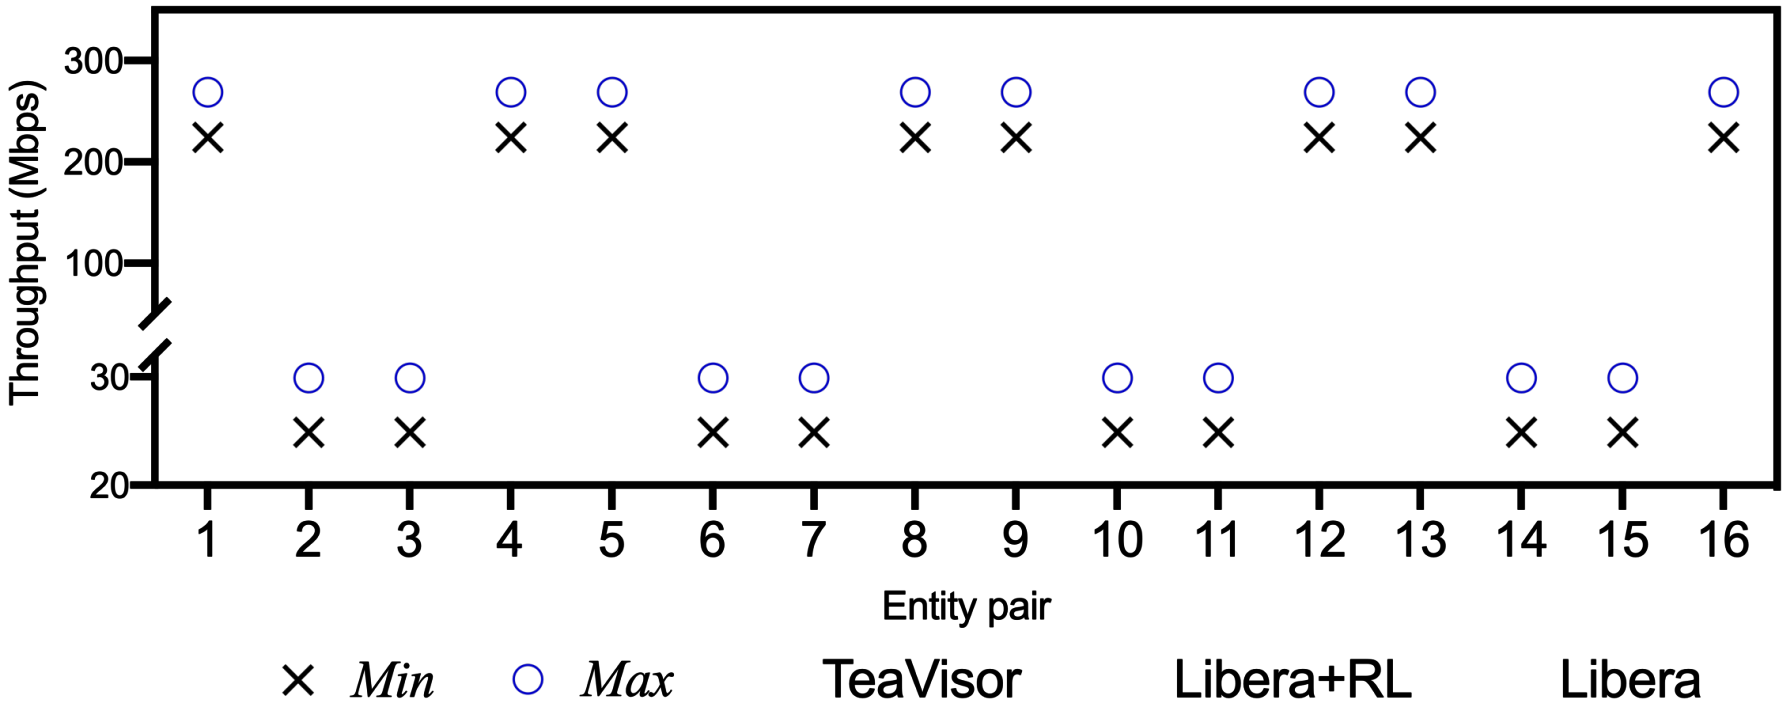

Bandwidth isolation guarantee - TeaVisor

- Zero error rates for bandwidth isolation guarantee
  - Always higher than *Min* and lower than Max

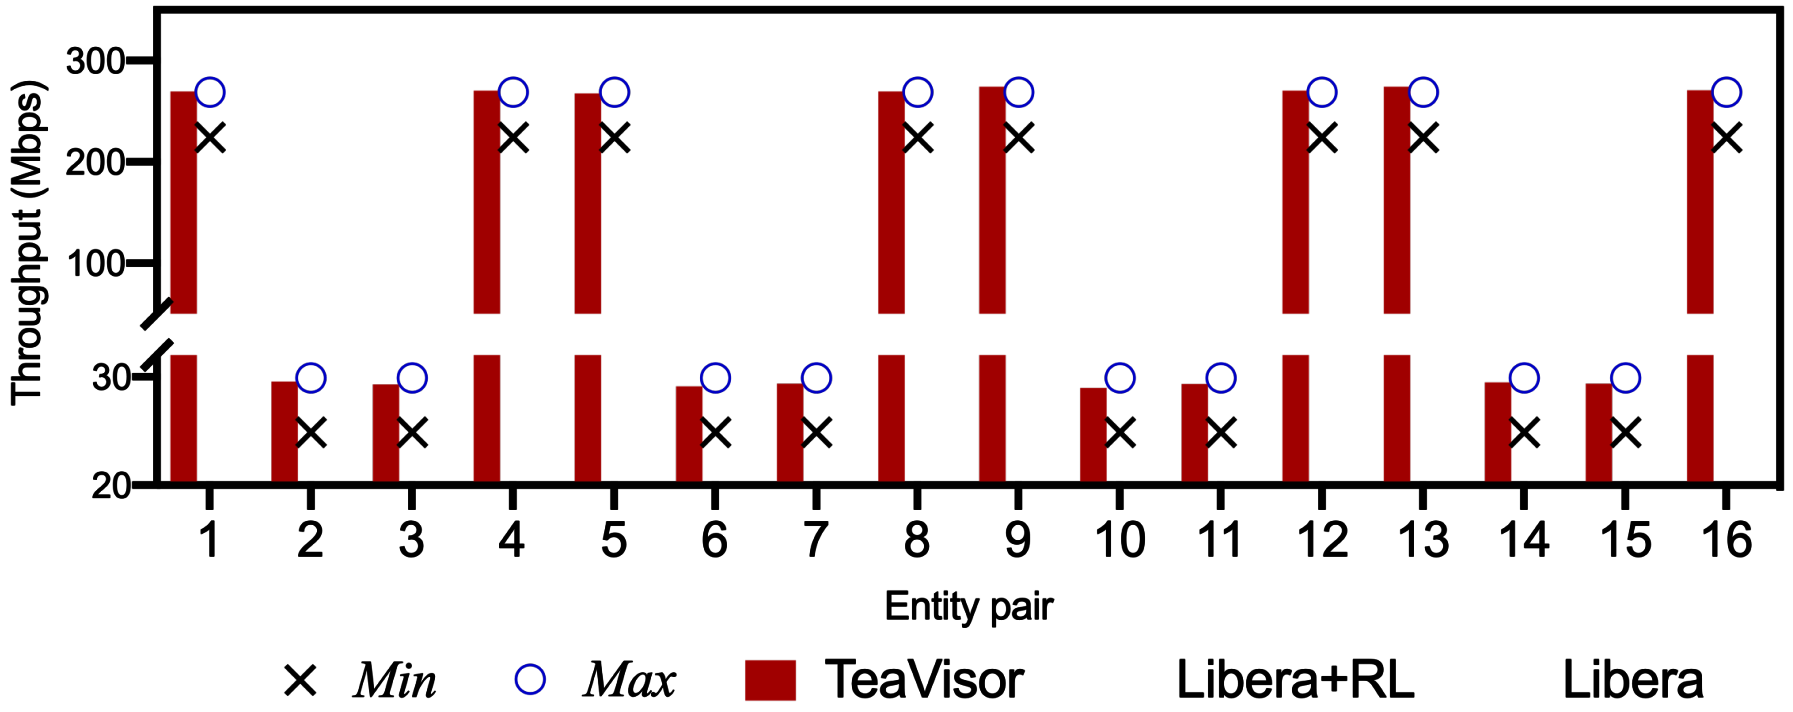

Bandwidth isolation guarantee – Libera+RL

- 42% less throughput than *Min* for 4, 5, 12, and 13 entity pairs
- No additional bandwidth even with the spare network capacity

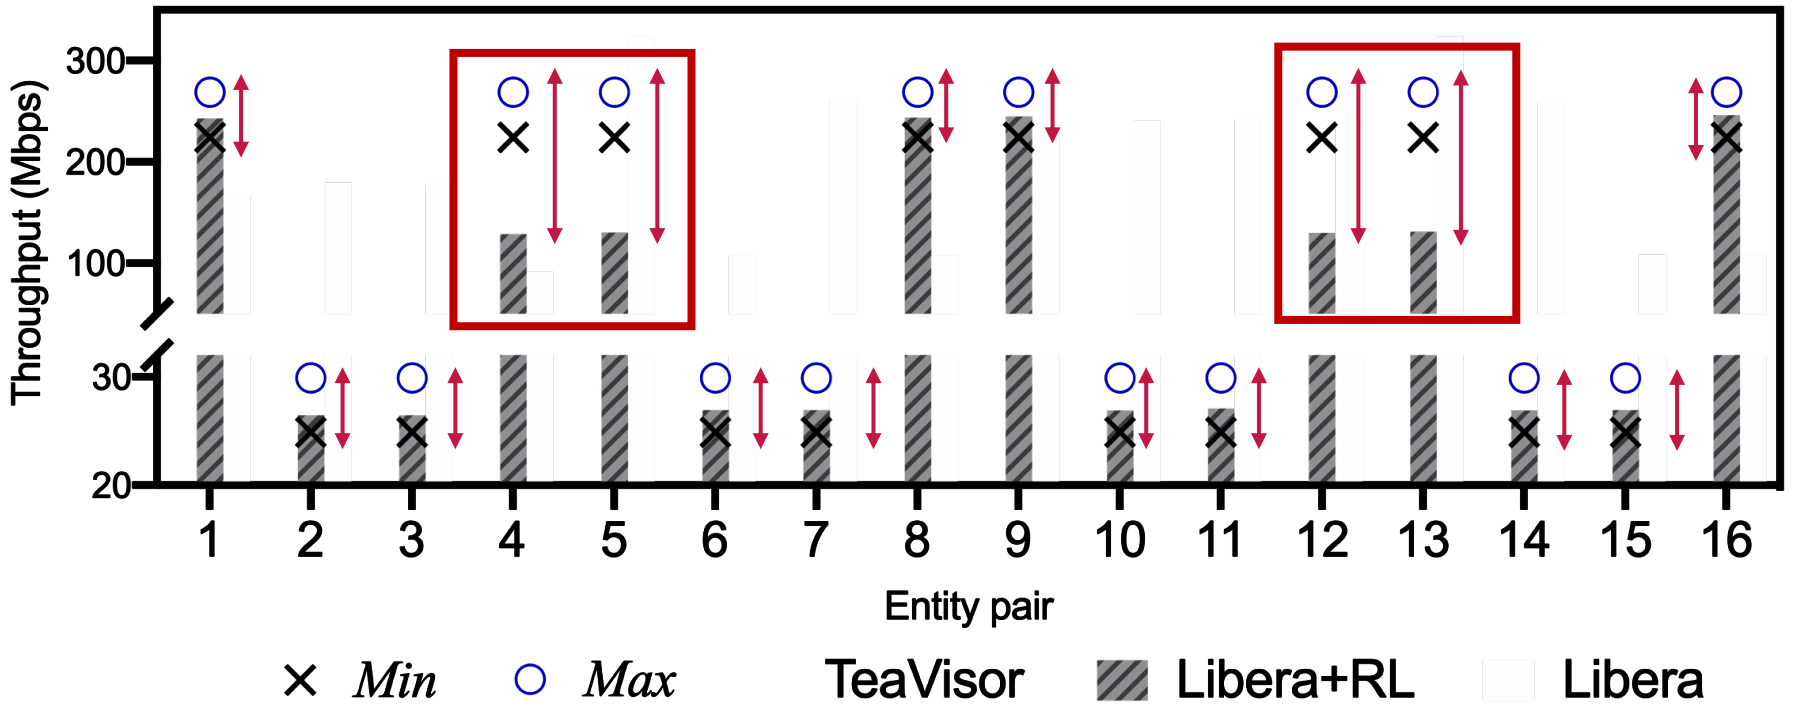

Bandwidth isolation guarantee – Libera

- 47% less throughput than *Min* for 1, 4, 8, and 16 entity pairs
- High bandwidth differences between entity pairs (lack of bandwidth management)

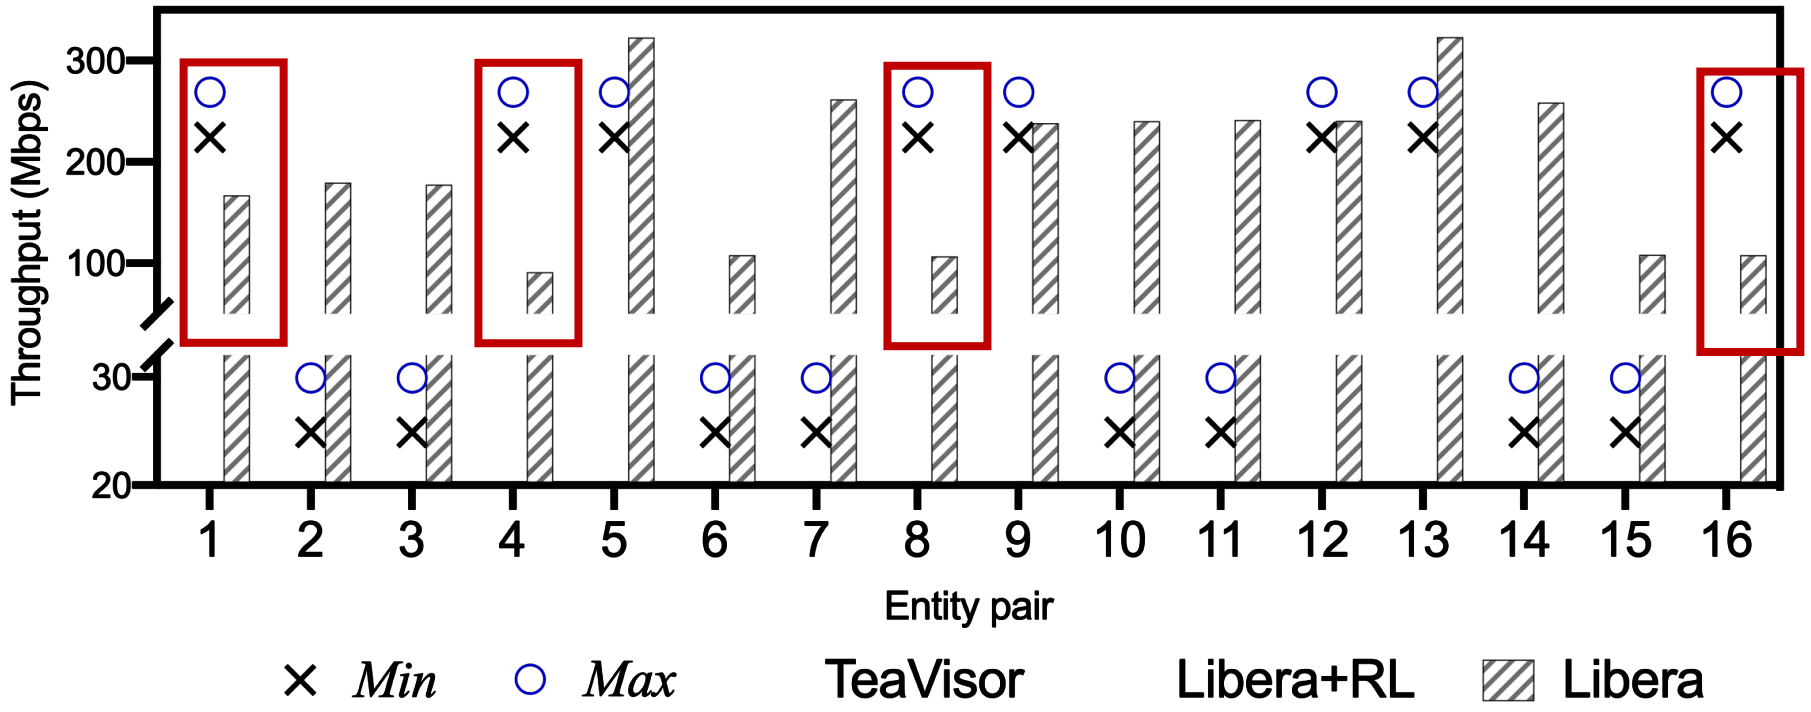

Bandwidth isolation guarantee

- TeaVisor: Near-zero error rates on both *Min* and *Max*
- Libera+RL: unsatisfaction on *Min*, always lower bandwidth than *Min* (poor utilization)
- Libera: unsatisfaction on *Min*, no restriction on *Max* requirement

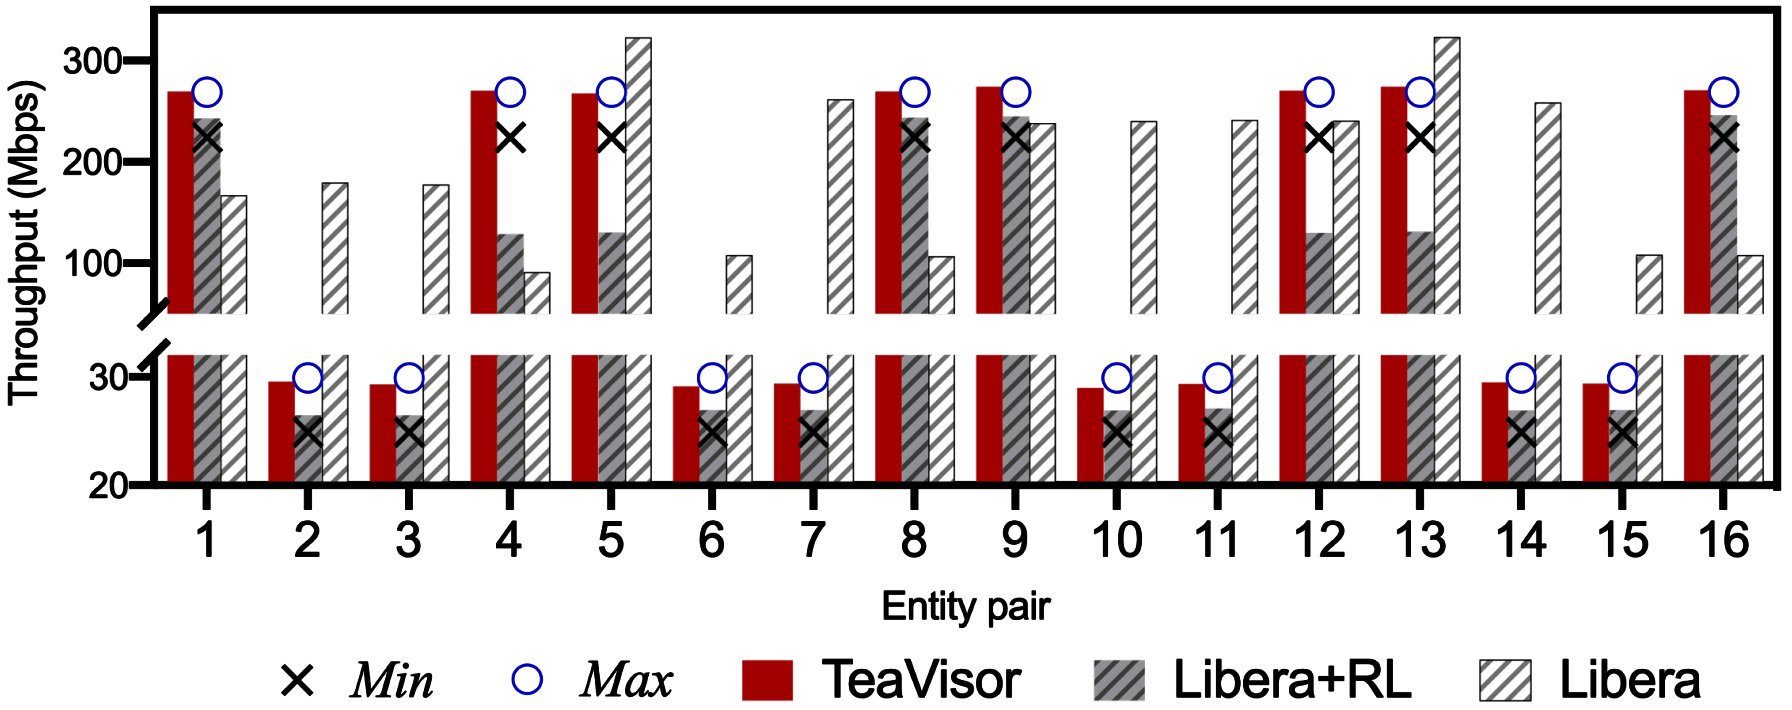

Bandwidth isolation guarantee – Error rate to *Min*

- Error rate to *Min/Max*
  - The extent to which bandwidth given to entity pair is higher or lower than *Min/Max*

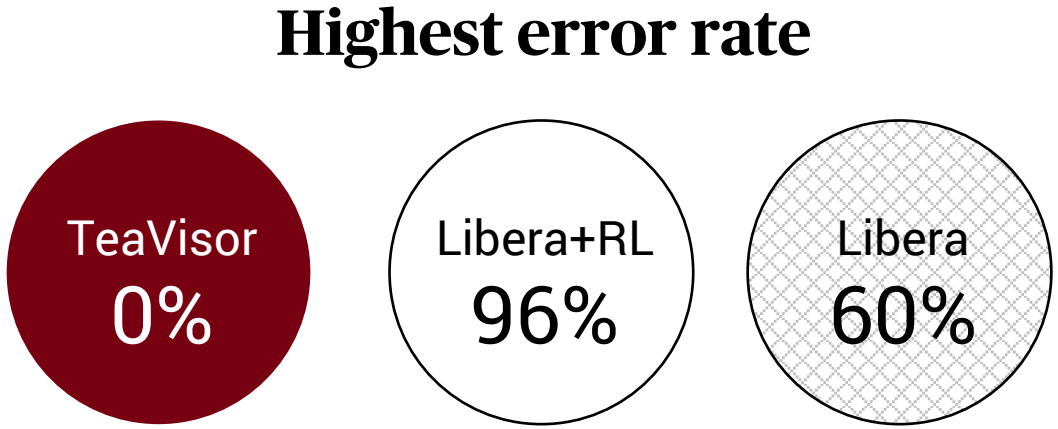

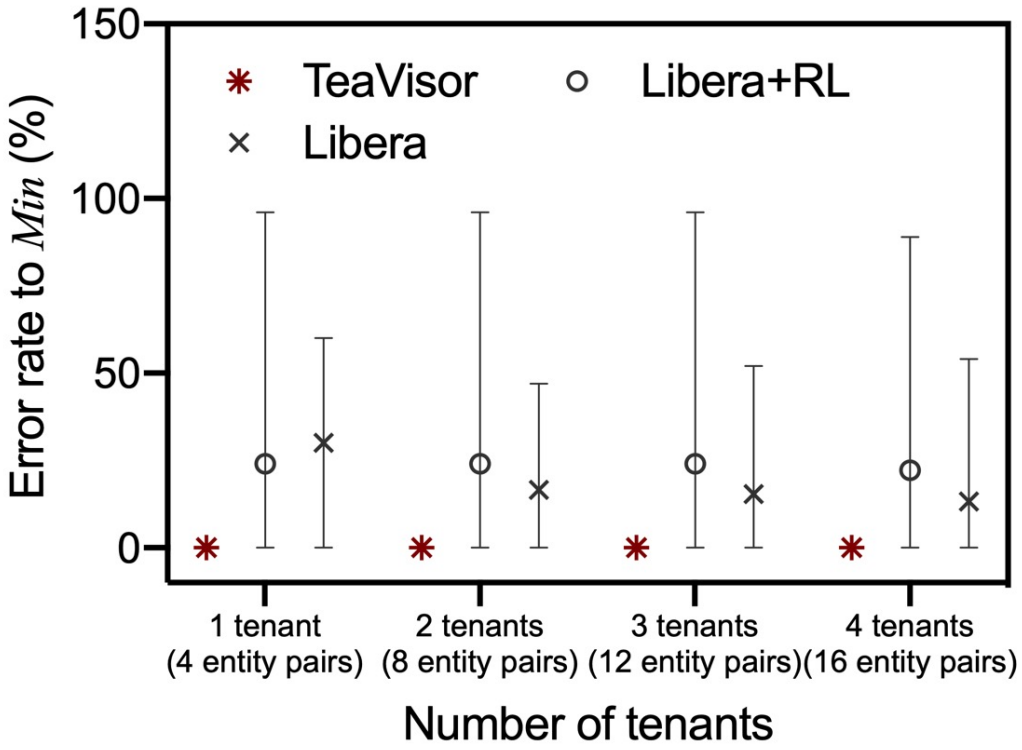

Bandwidth isolation guarantee – Error rate to *Max*

Highest error rate

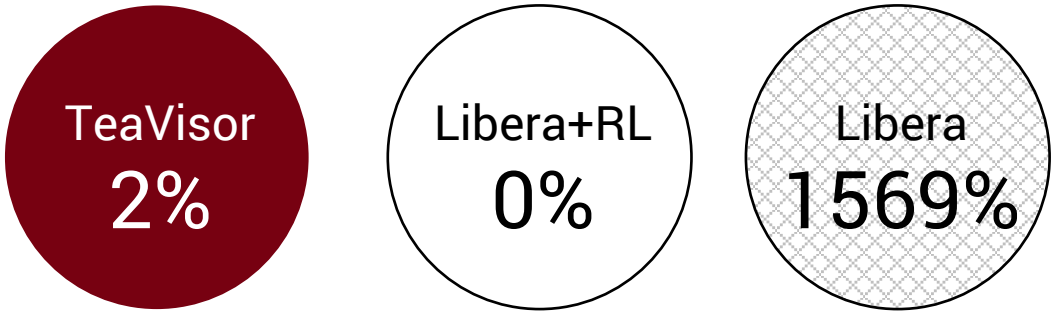

- Libera+RL always provide bandwidth lower than *Min*

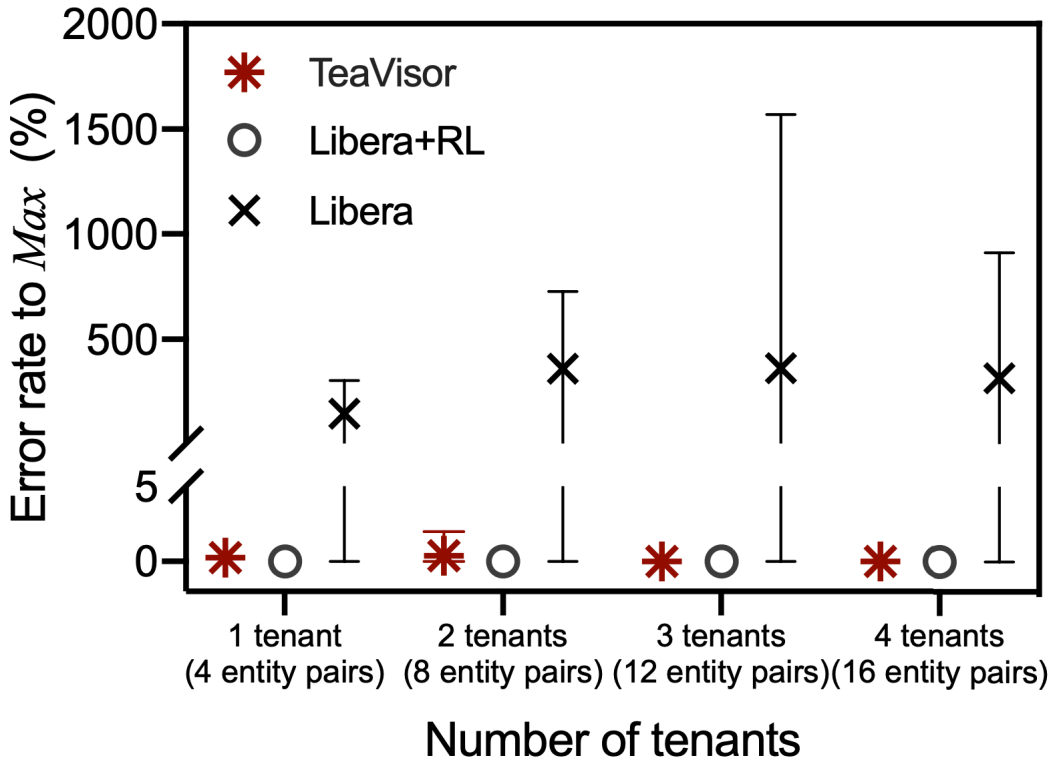

Overheads of TeaVisor

- Control traffic consumption: 2.6 and 2.9 times higher than Libera+RL and Libera
- Memory consumption: similar to others (99.7% of Libera)

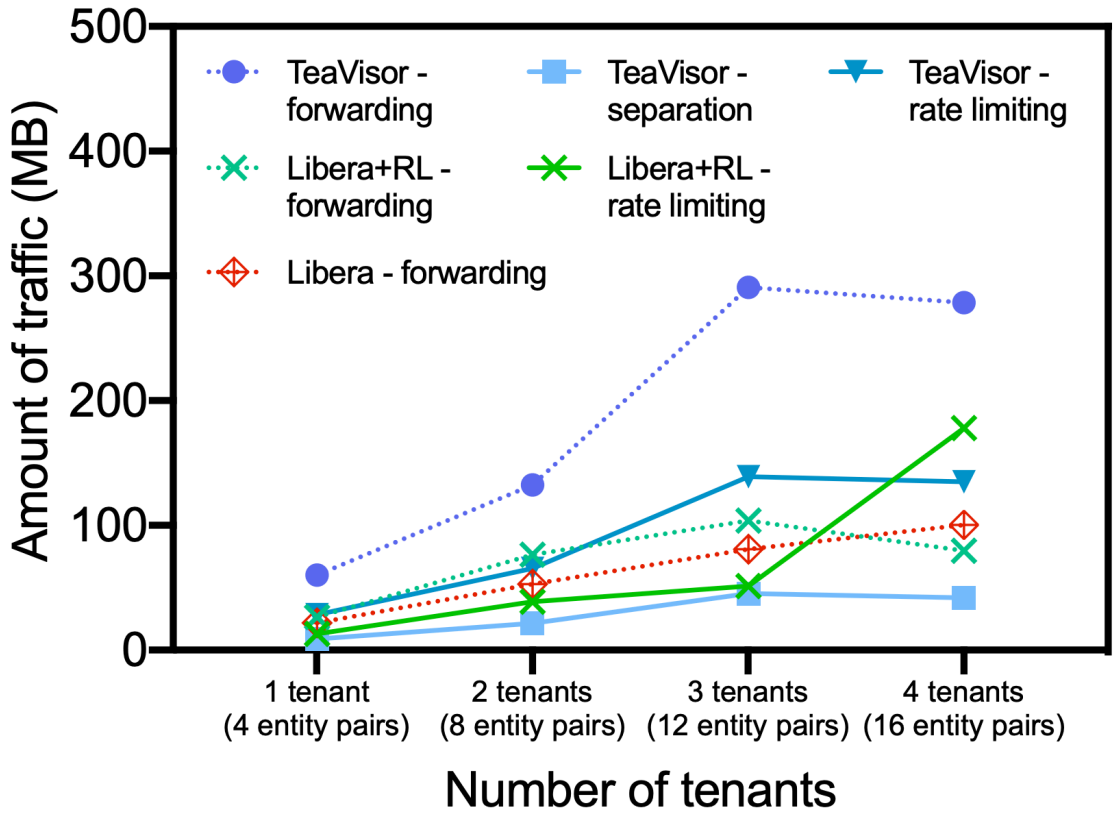

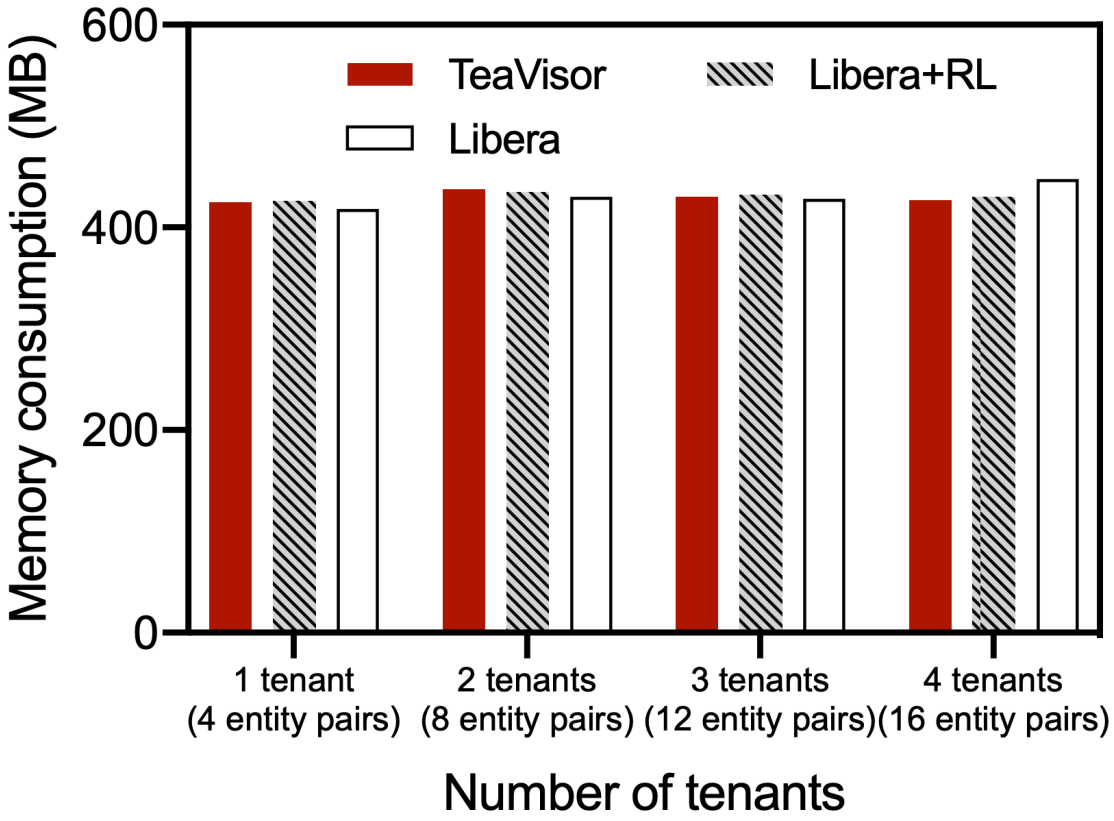

Conclusion

- Bandwidth isolation guarantee for predictable and stable services in clouds
- TeaVisor: the first bandwidth isolation guarantee framework for SDN-NV
  - Path virtualization, bandwidth reservation, and path establishment
- Near-zero error rates for both minimum and maximum bandwidth requirements

THANK YOU!

Gyeongsik Yang

ksyang@os.korea.ac.kr | gsyang33.github.io

KOREA UNIVERSITY