Network-Scale Virtual Output Queues

Brian Petersen, Ethernovia, Inc.

IEEE SA Ethernet & IP Automotive Technology Day

Yokohama, Japan—November 2022

ETHERNOVIA

Agenda

- IVN QoS
- Causes of Latency and Packet Loss
- Queues and Networks
- Intro to VOQs
- VOQ Behavior Analysis
- Shaping Algorithms
- Conclusions

In-Vehicle Network QoS Objectives

- Bounded latency
- Minimal packet loss

What causes latency?

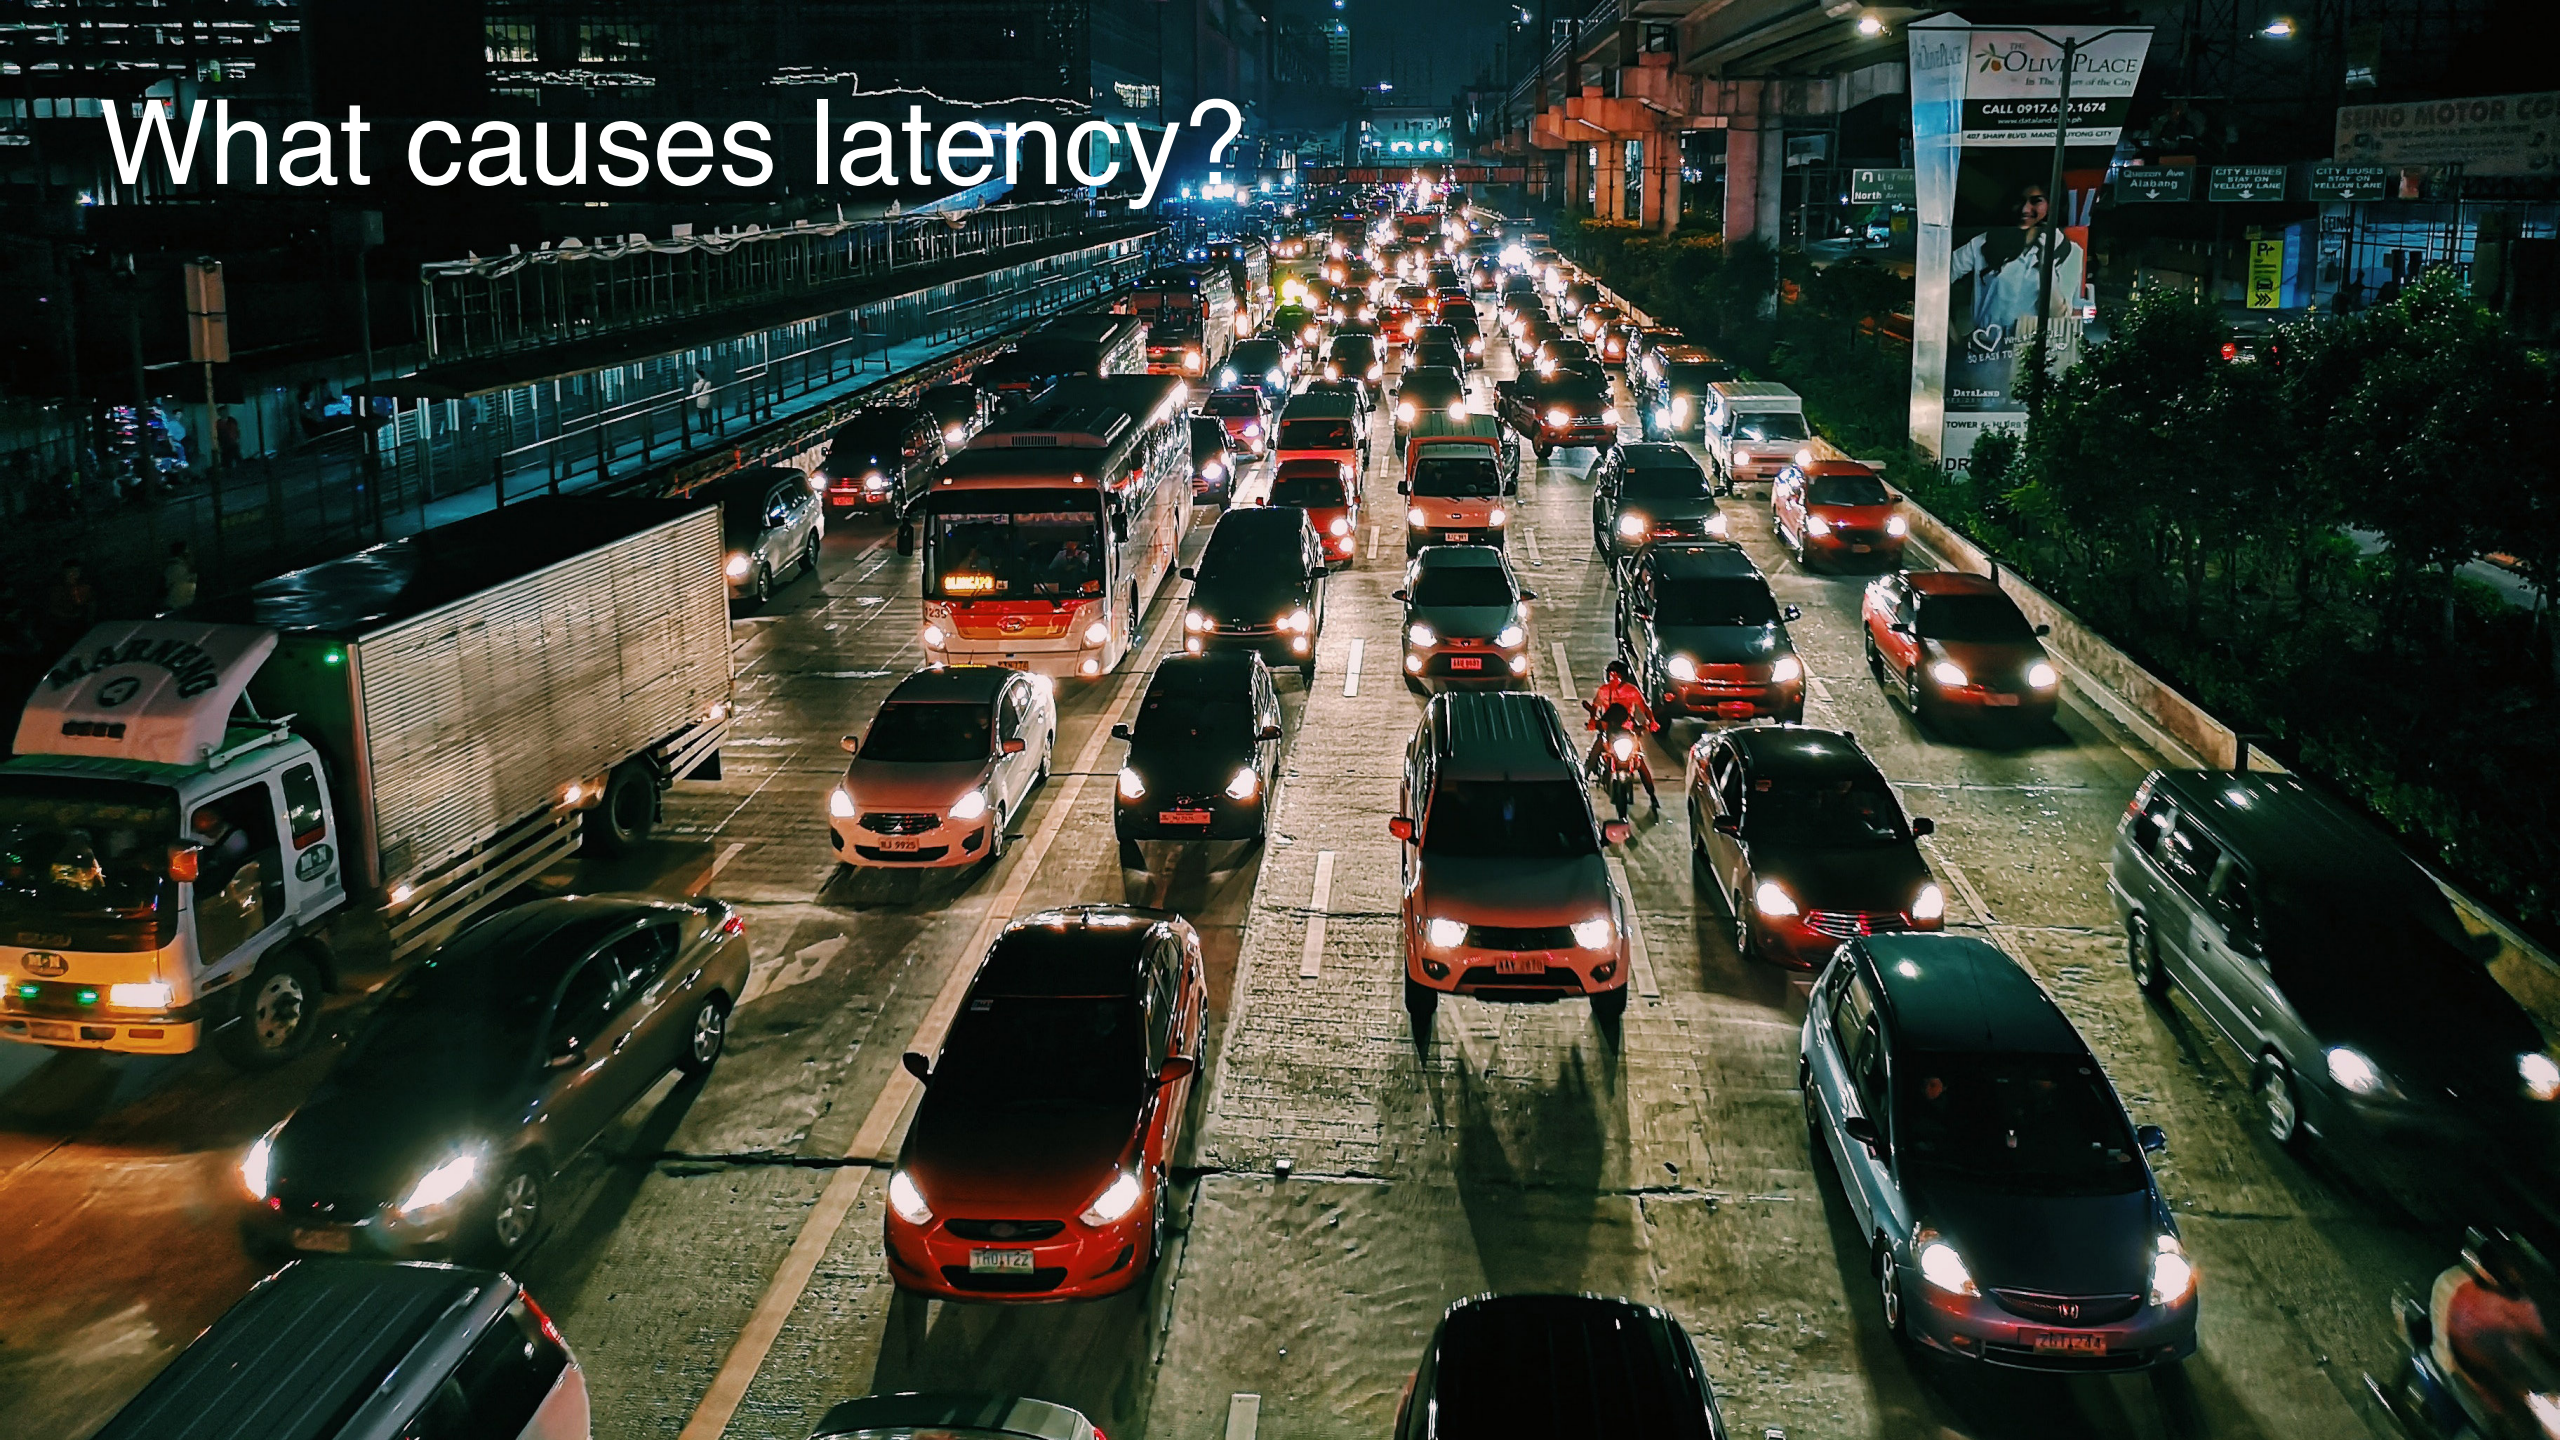

What Causes Packet Loss?

- Bit errors
  - 10^{-12} BER—1/12,000,000 lost packets
- Buffer overruns
  - 10:1 rate disparity—9/10 lost packets

Congestion is the Real Problem

- Increases queue occupancy
- Increases risk of packet loss

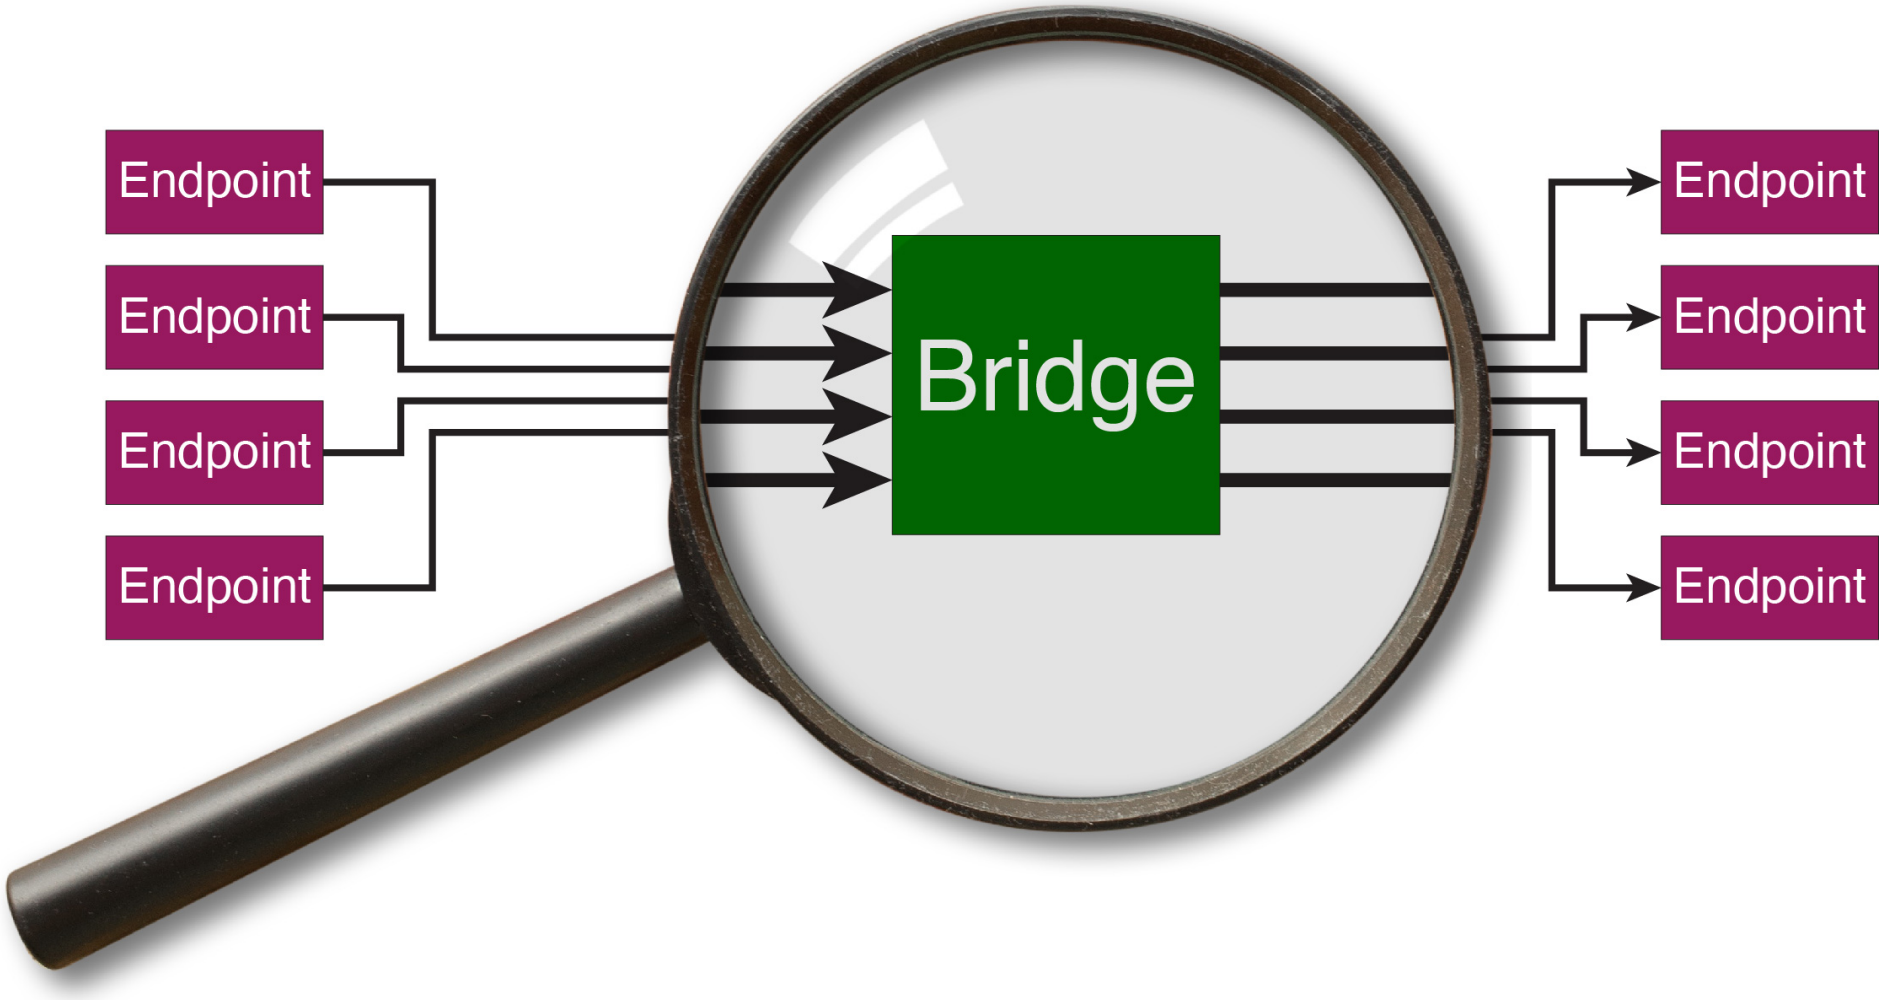
Endpoint
Endpoint
Endpoint
Endpoint
Bridge
Endpoint
Endpoint
Endpoint
Endpoint

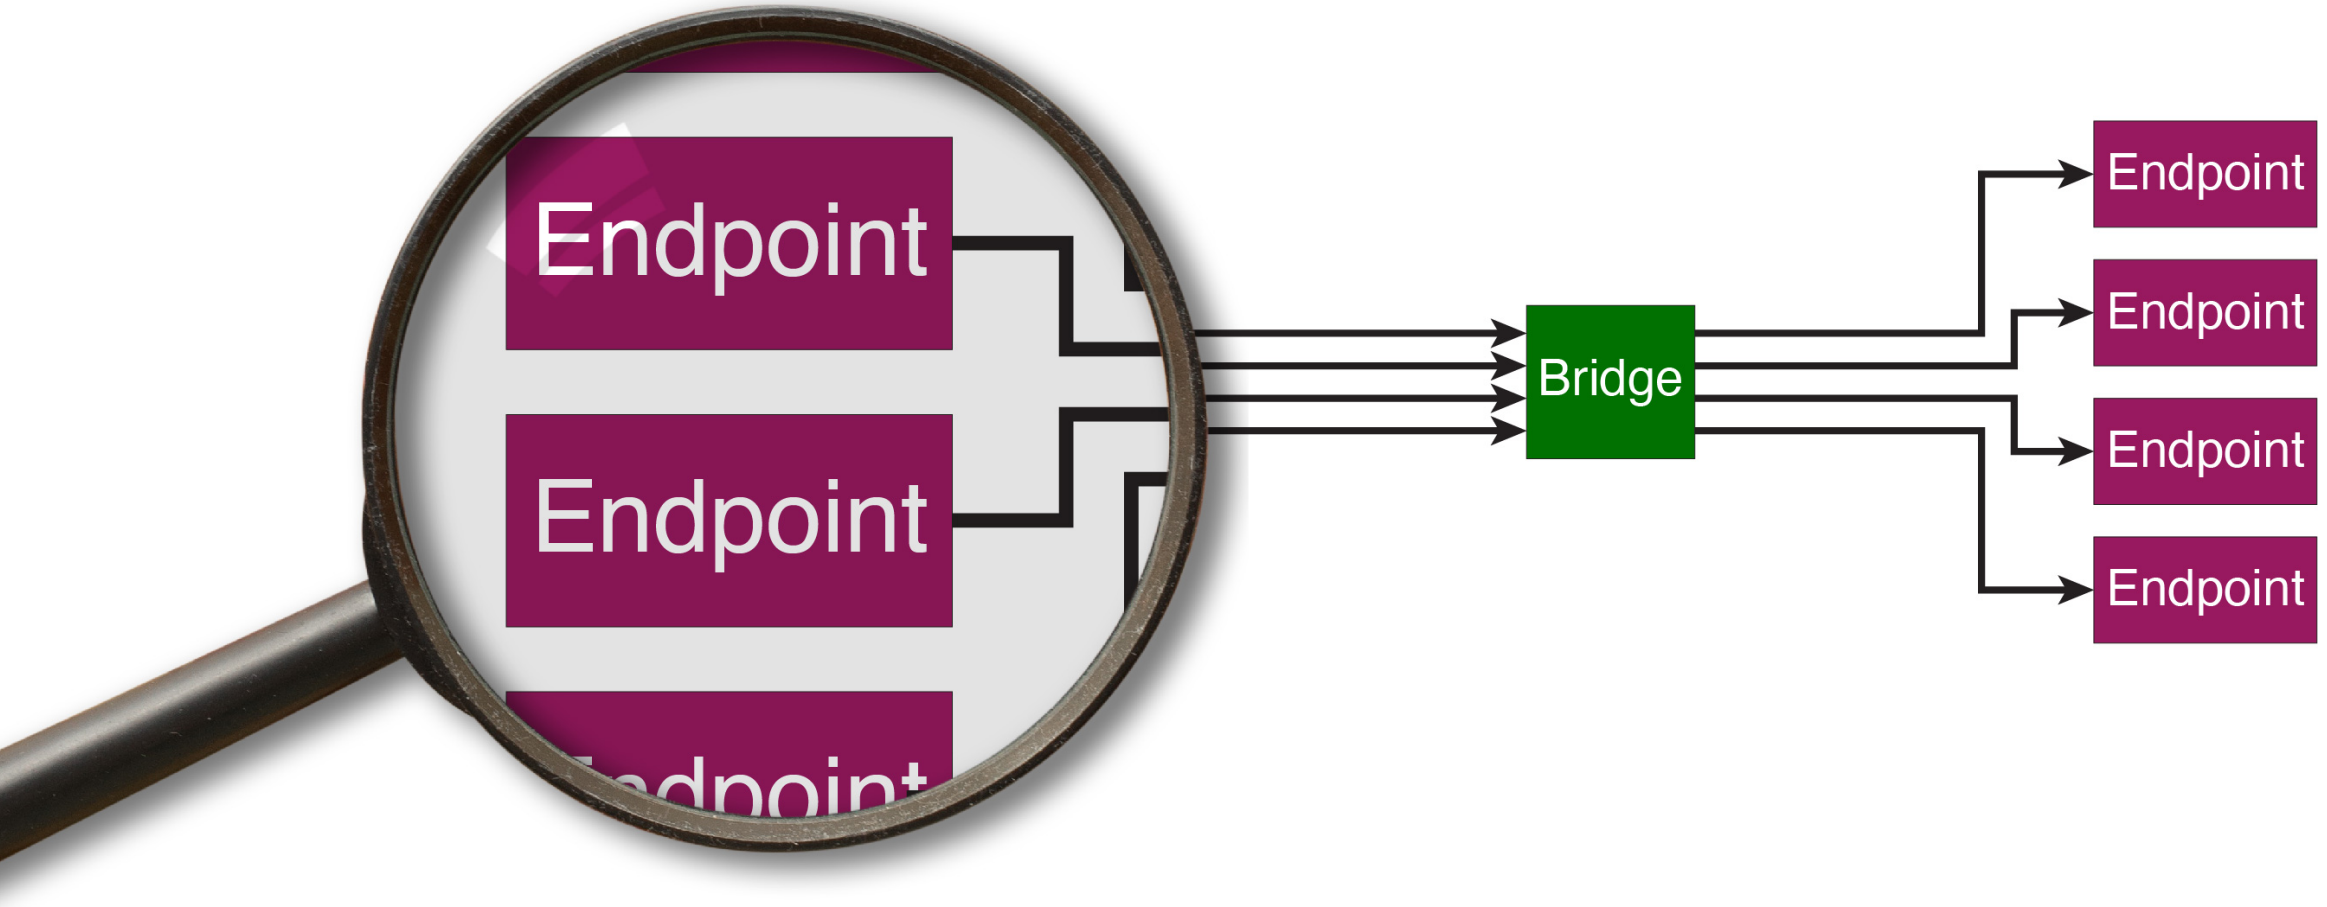
Endpoint
Endpoint
Bridge
Endpoint
Endpoint
Endpoint
Endpoint

A Network of Queues

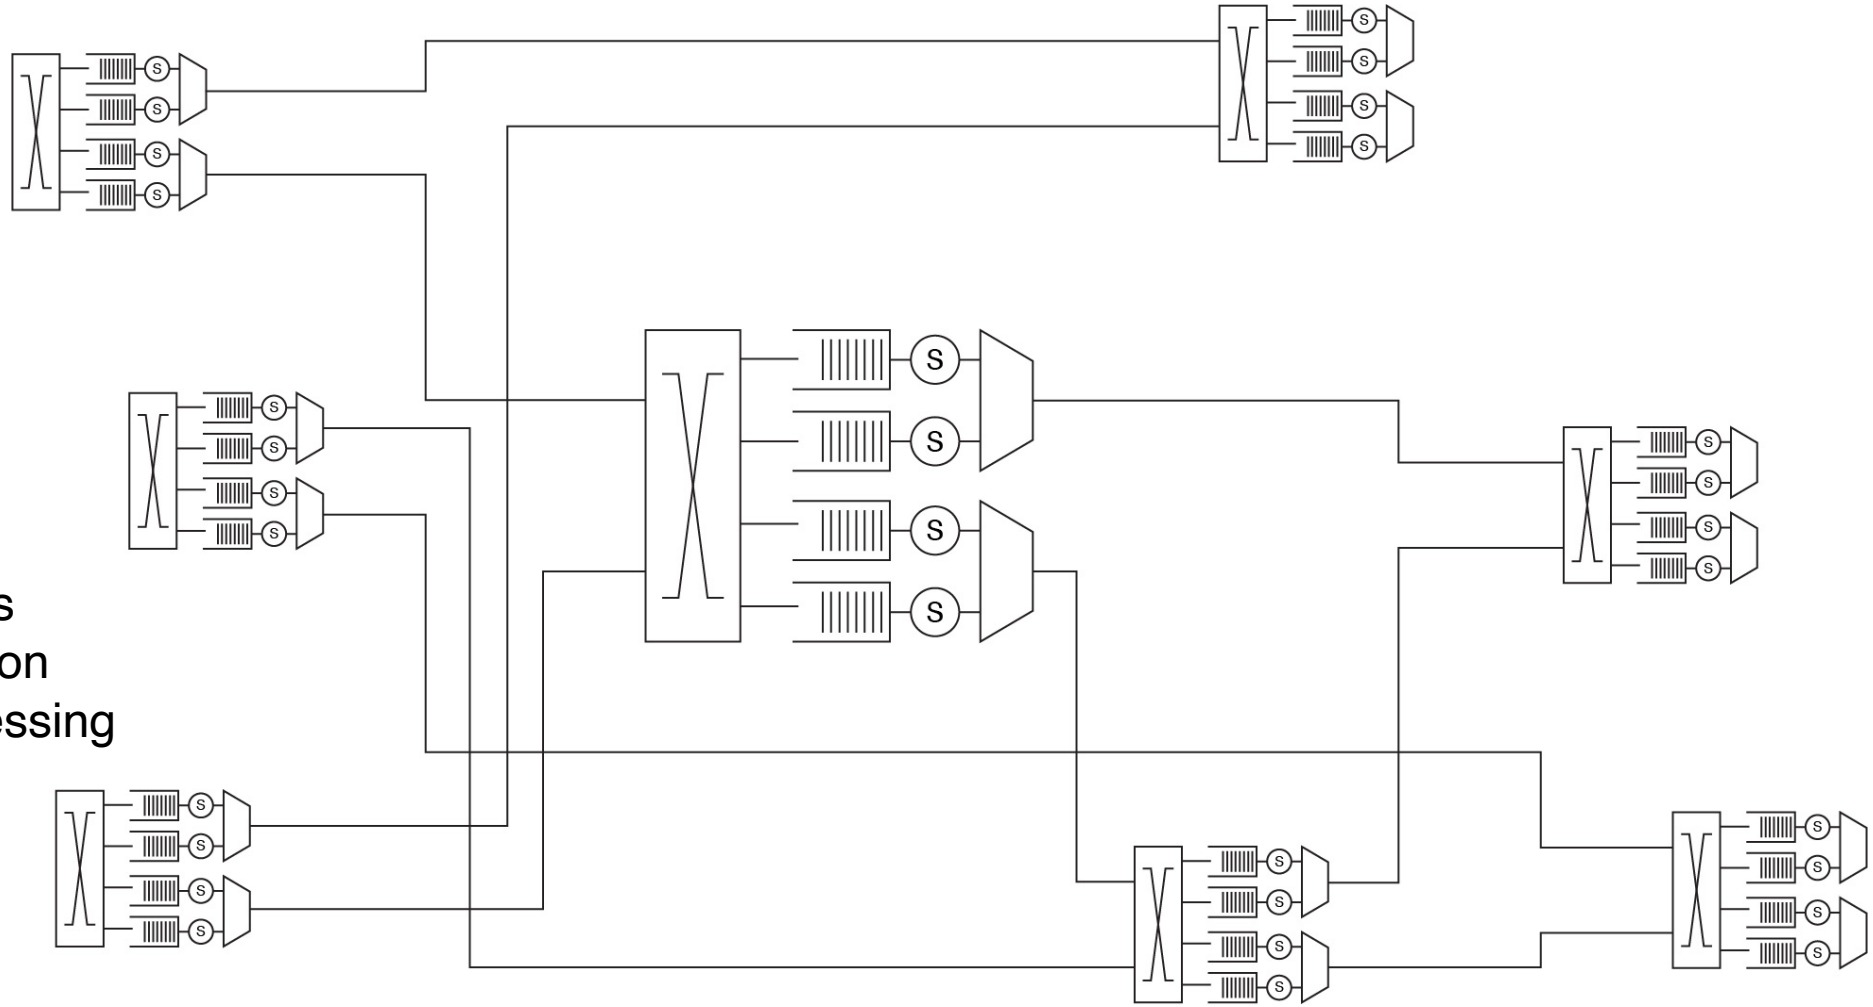

- servo commands
- display information
- distributed processing

Rate Limits Prevent Congestion

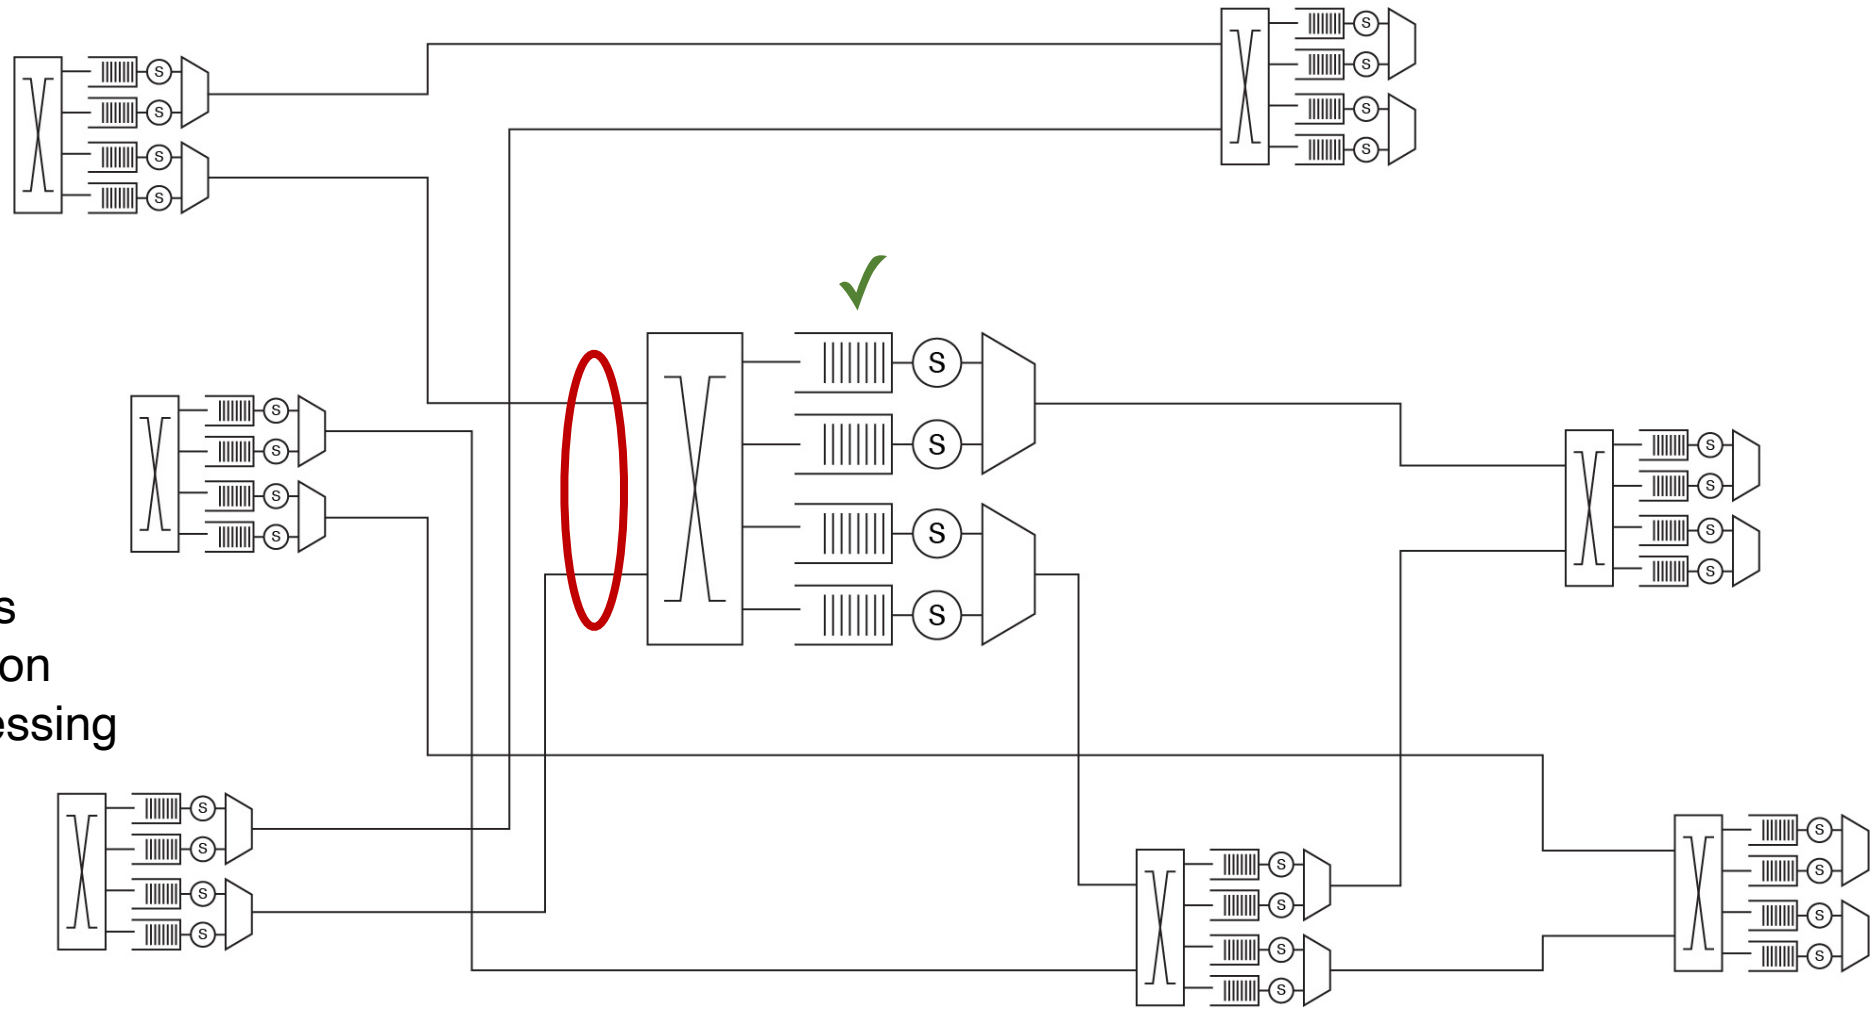

- servo commands
- display information
- distributed processing

Sources of Traffic

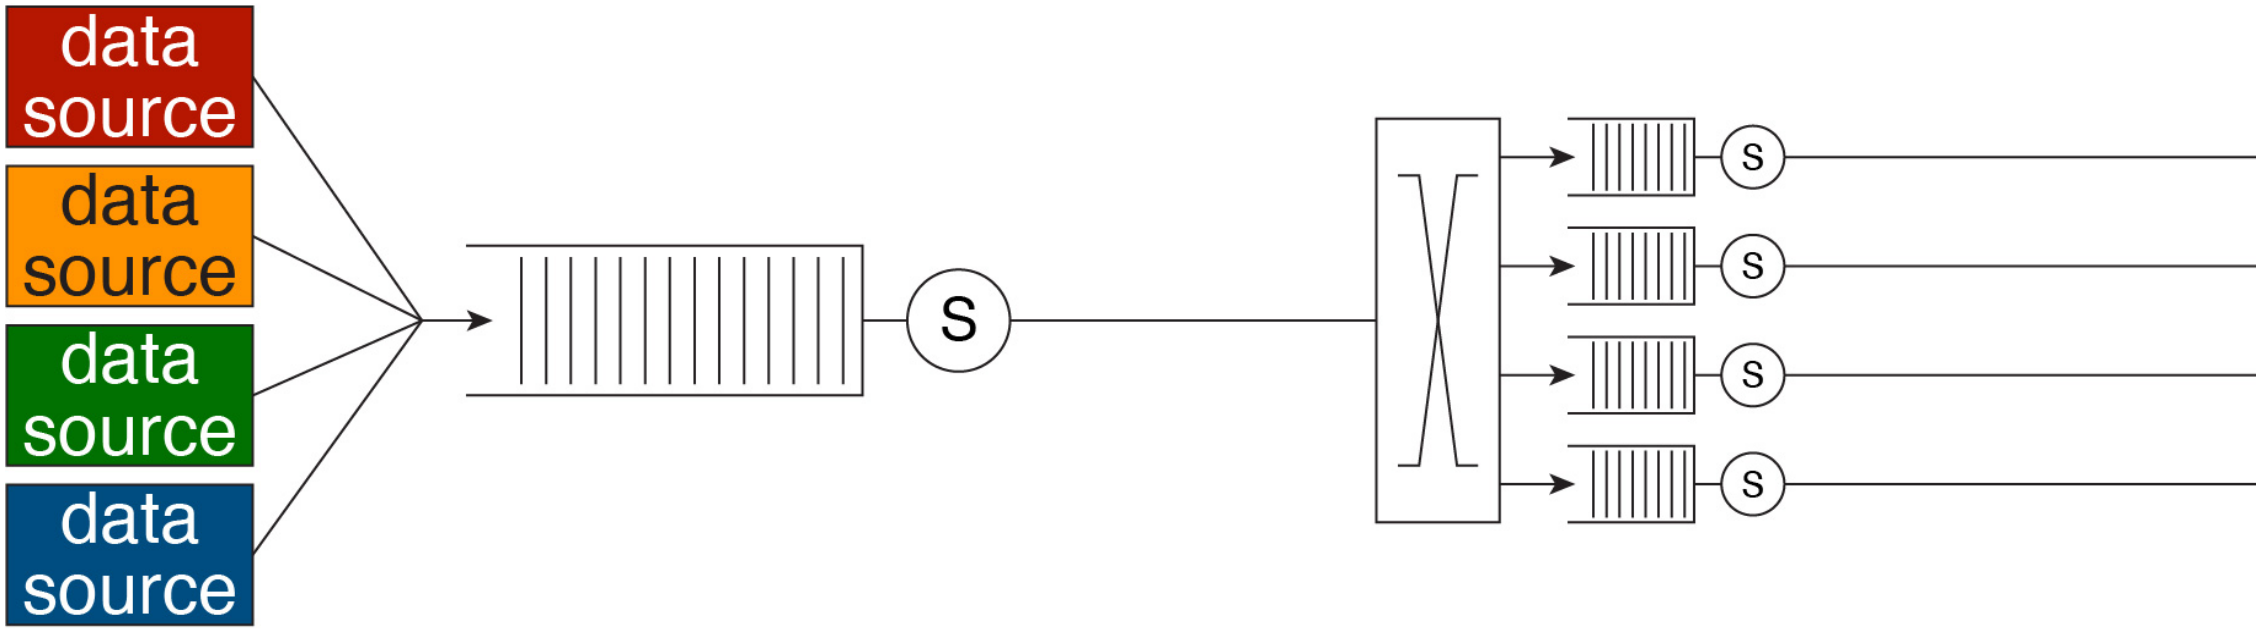

Head-of-Line Blocking

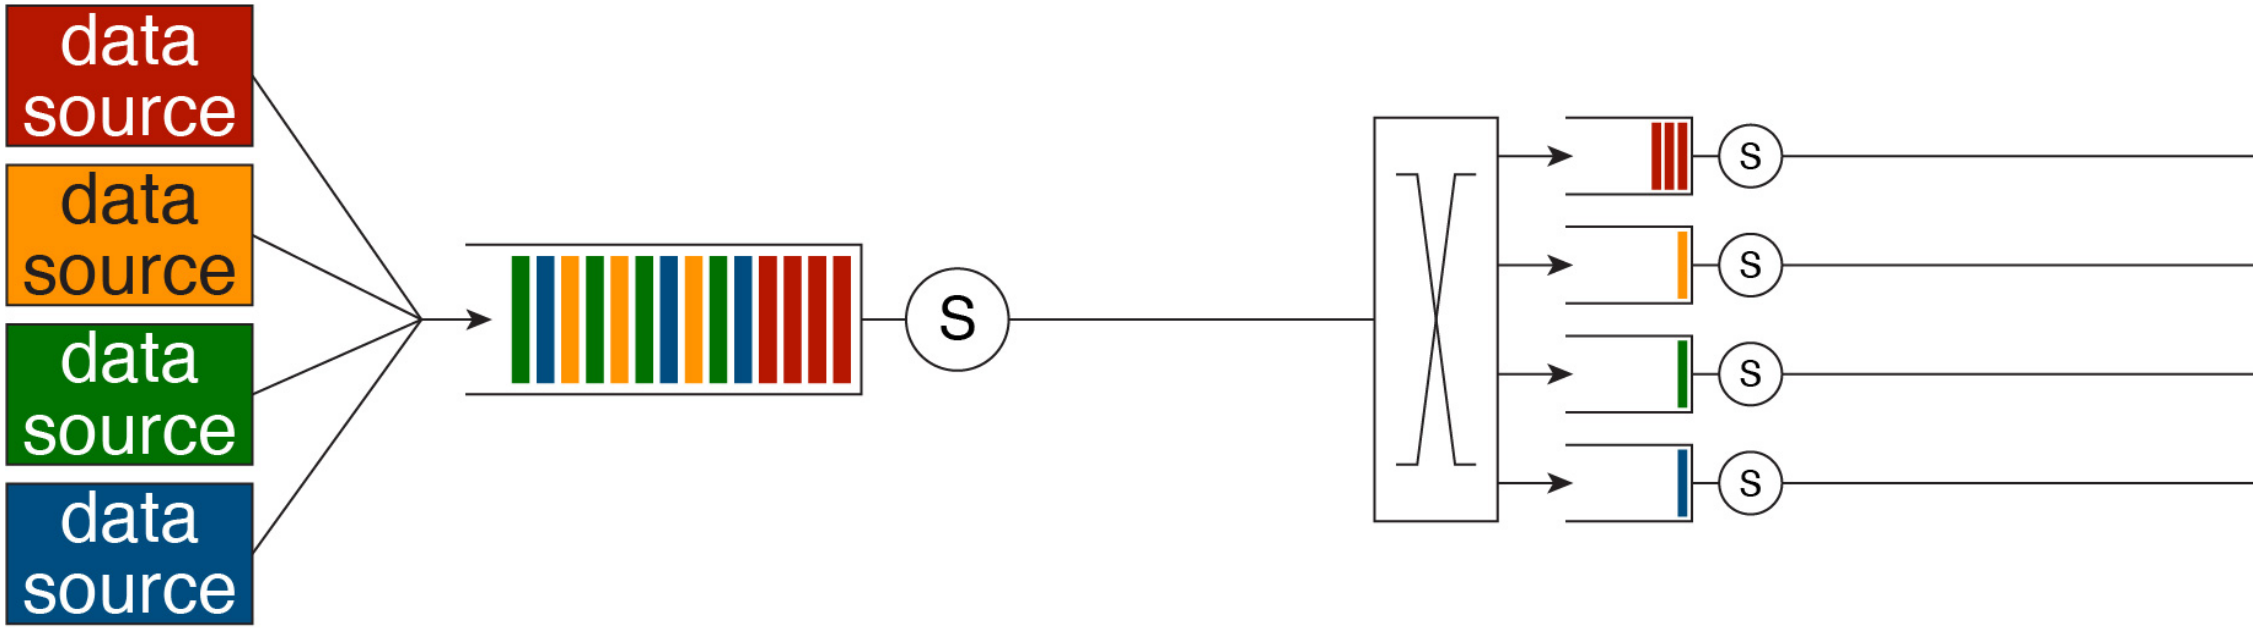

Virtual Output Queues

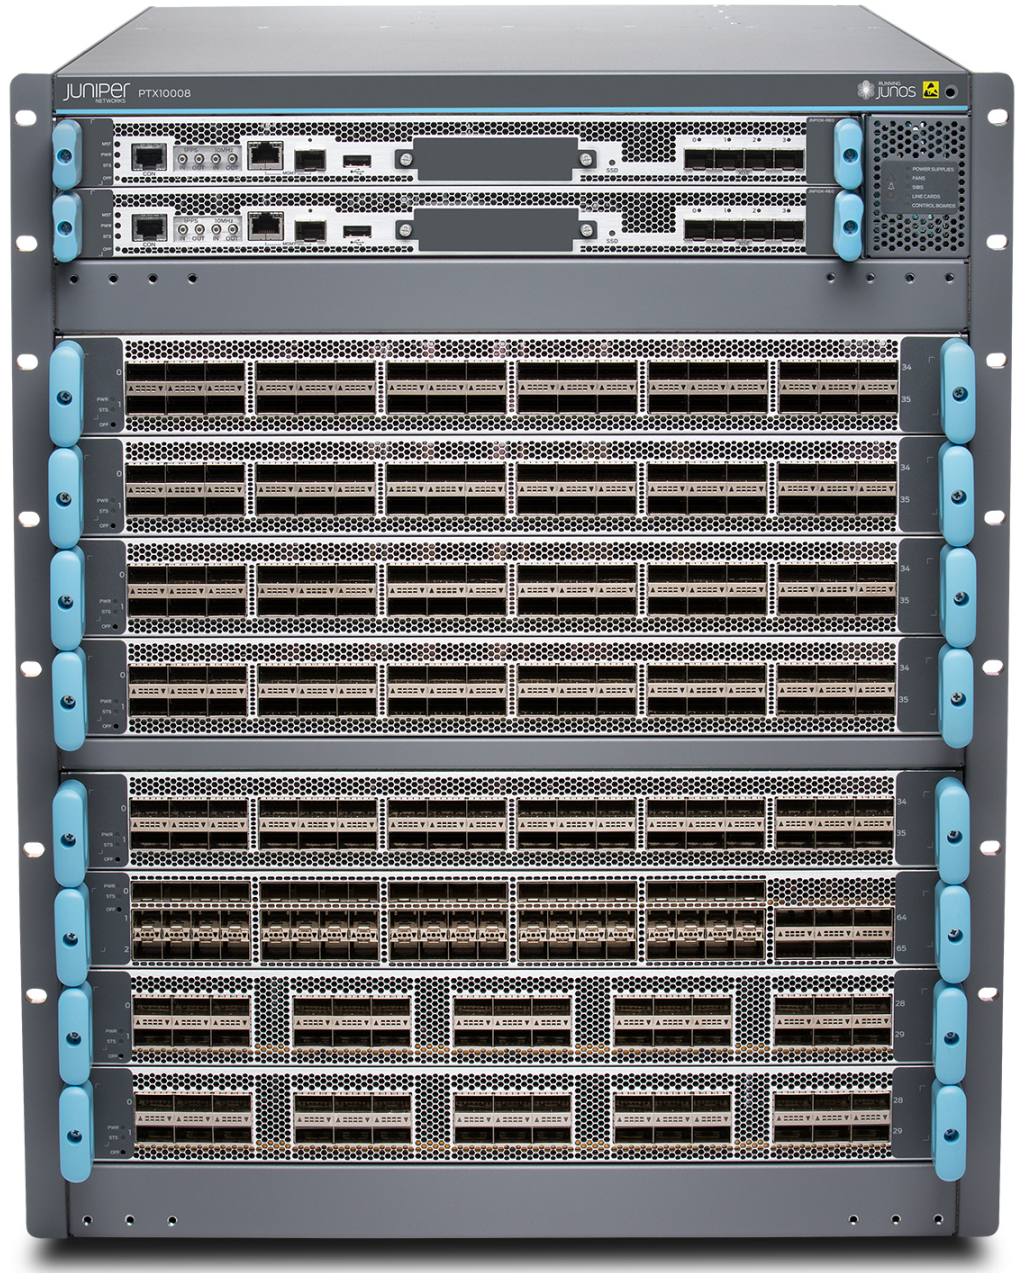

ETHERNOVIA

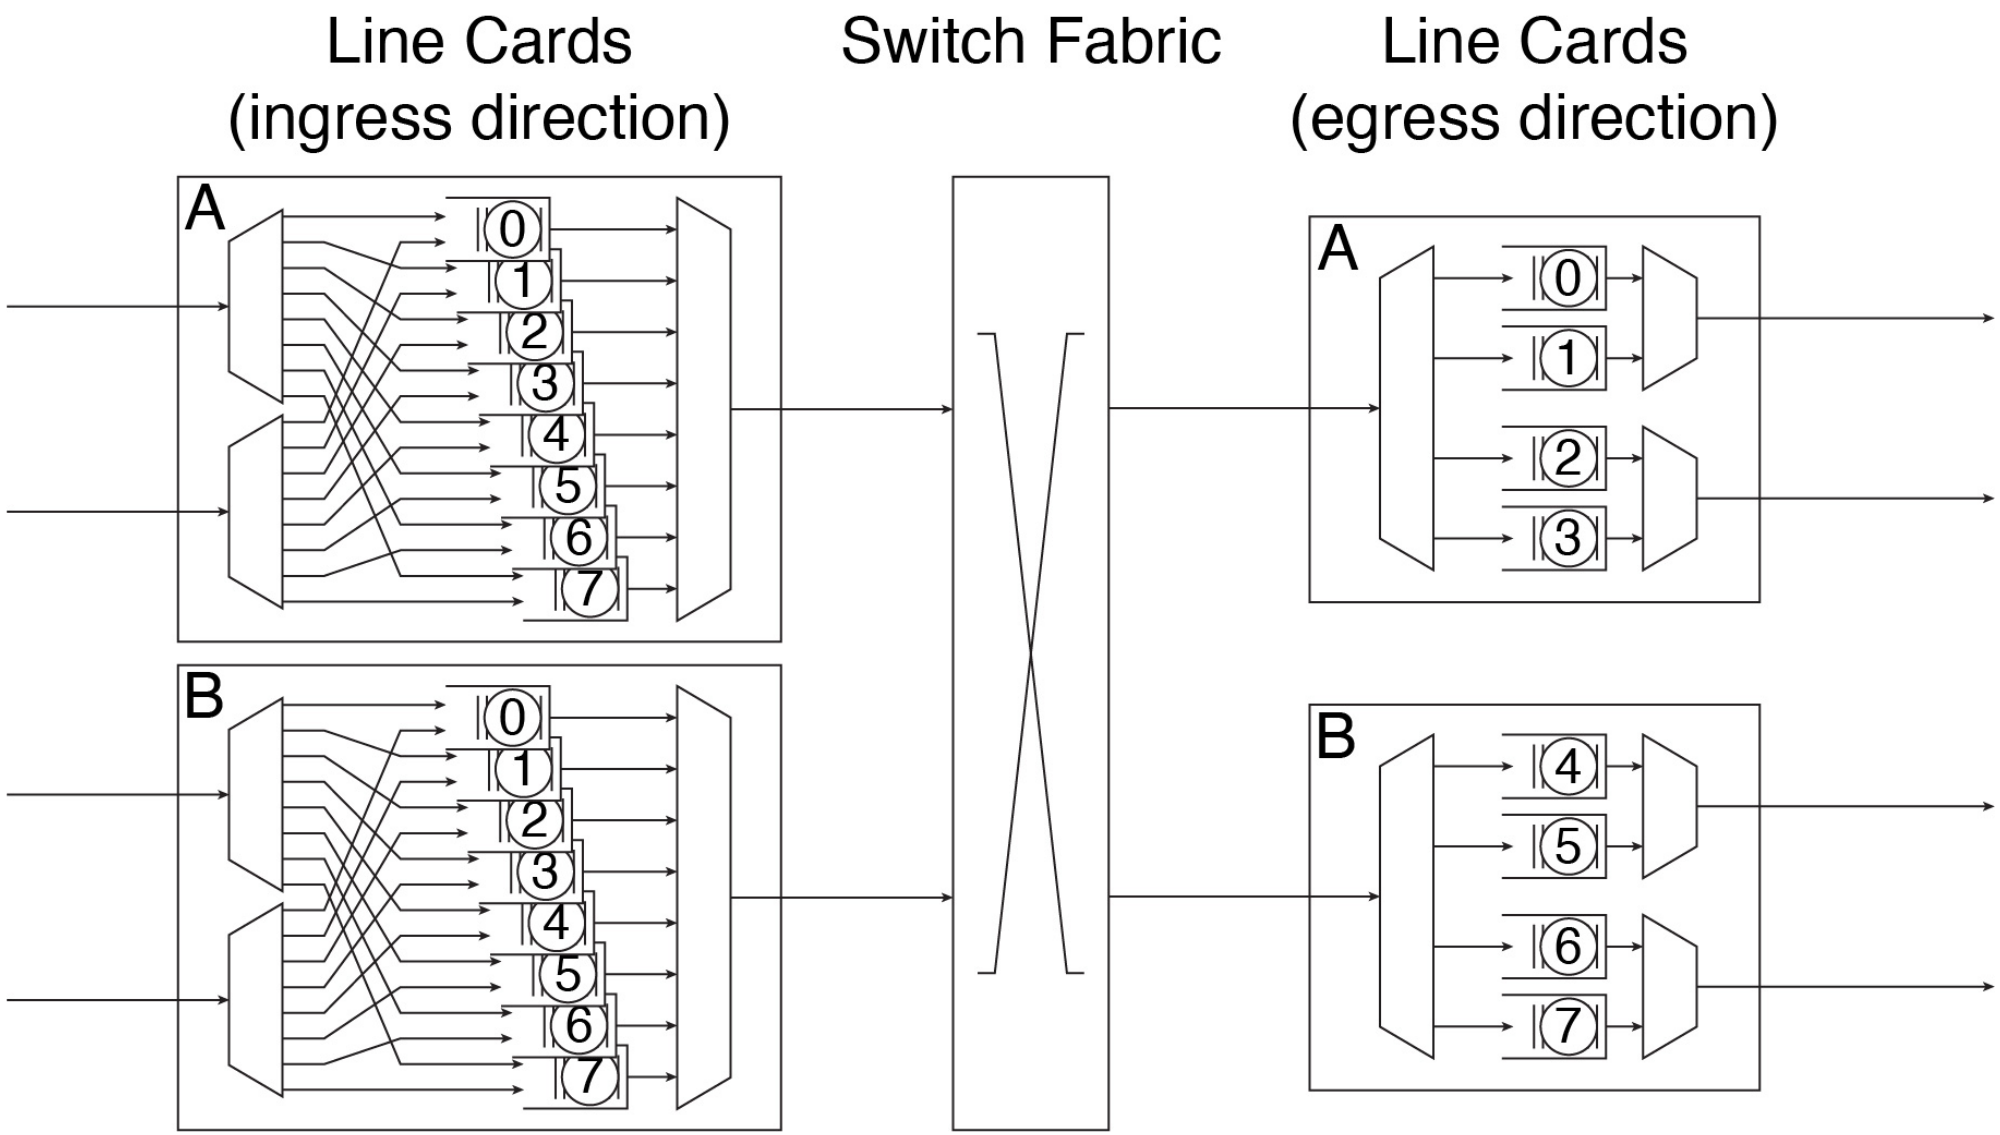
Line Cards
(ingress direction)
Switch Fabric
Line Cards
(egress direction)
A
0
1
2
3
4
5
6
7
B
0
1
2
3
4
5
6
7
A
0
1
2
3
B
4
5
6
7

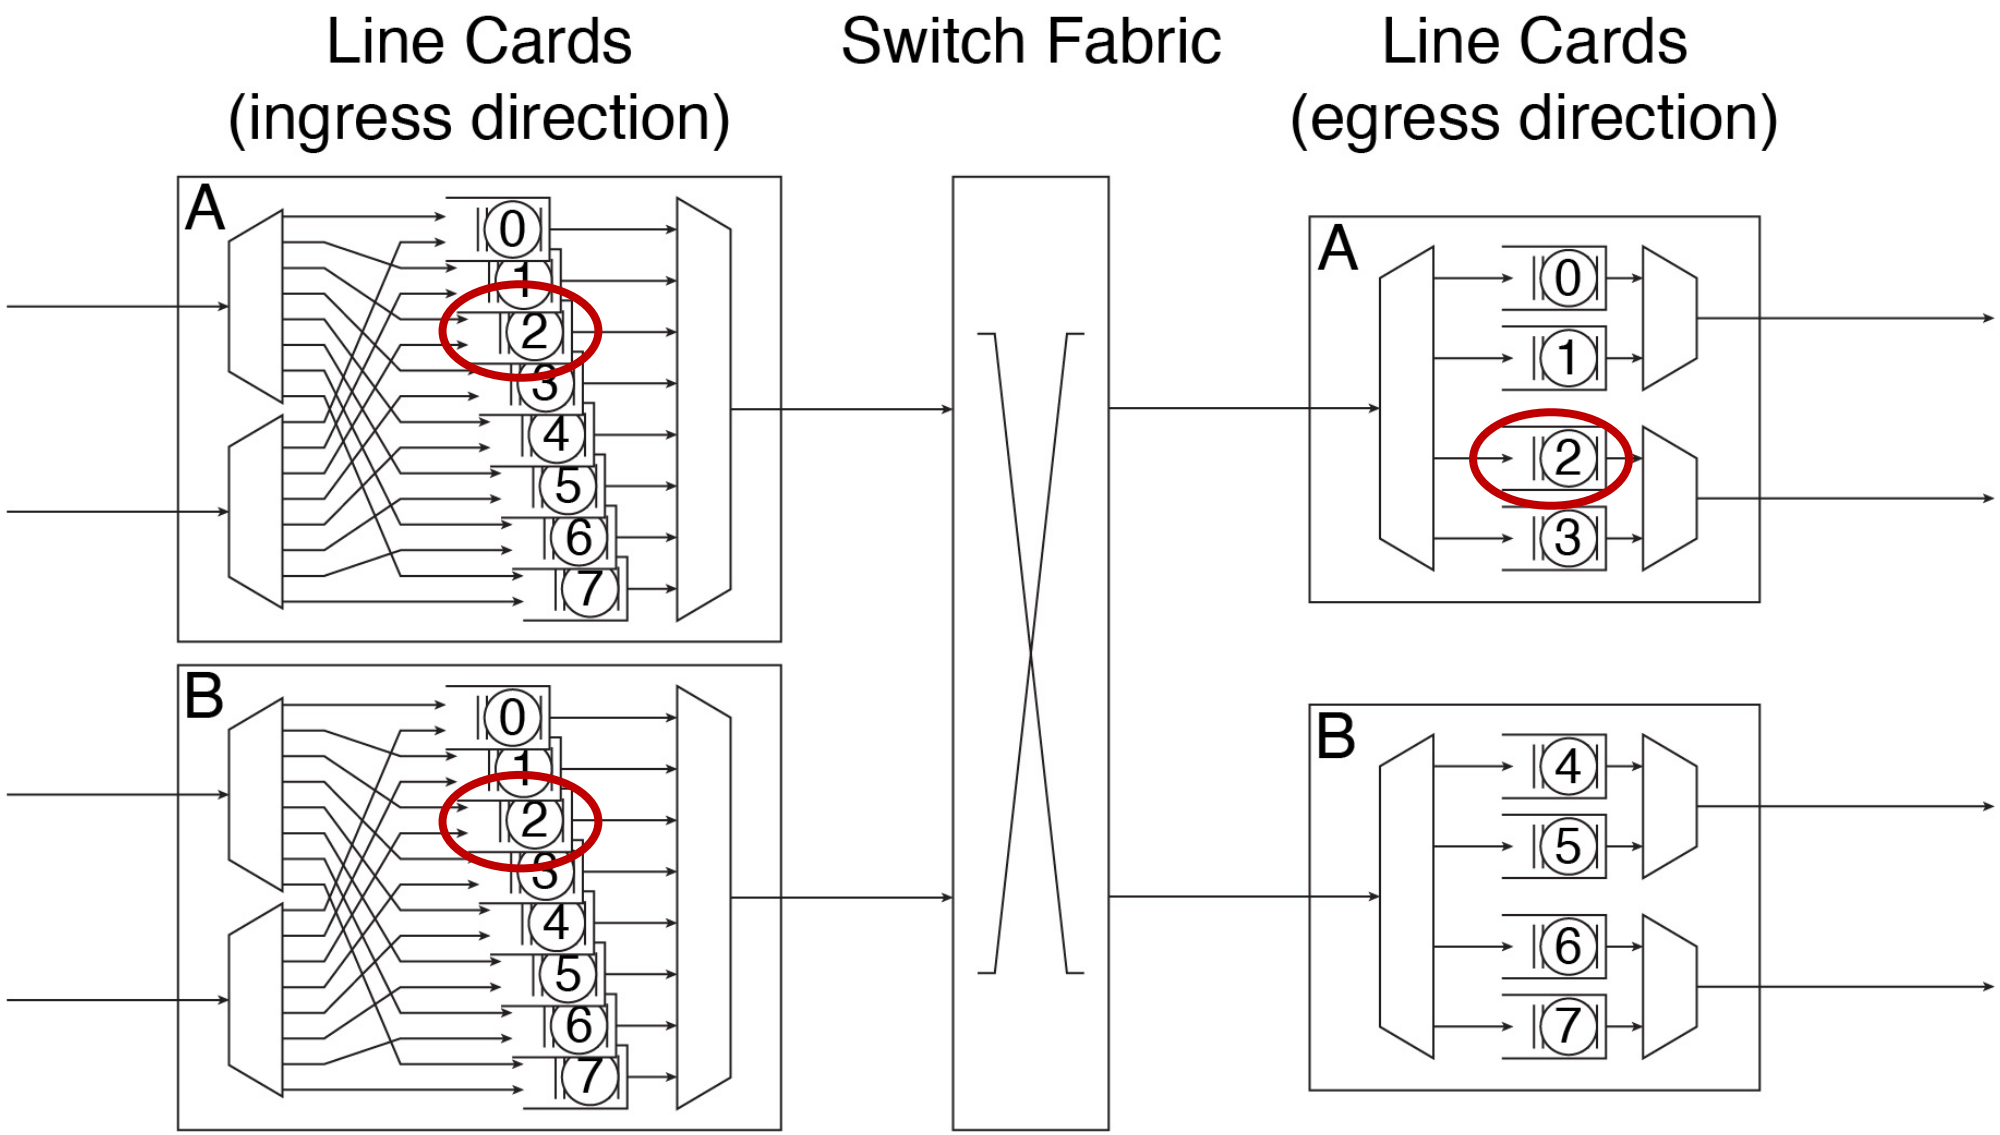
Line Cards
(ingress direction)
Switch Fabric
Line Cards
(egress direction)
A
0
1
2
3
4
5
6
7
B
0
1
2
3
4
5
6
7
A
0
1
2
3
B
4
5
6
7

More About Chassis Fabrics

- Non-blocking
  - Clos topology
- Single priority
  - priorities handled by ingress VOQs
- Shallow buffers
  - well-behaved ingress

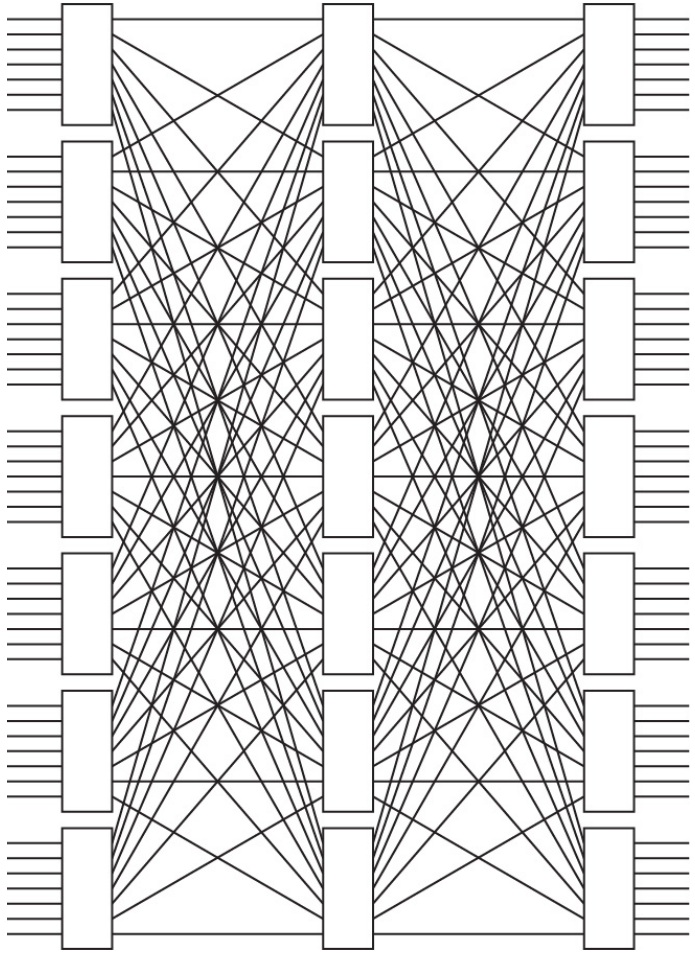

Applying Chassis Concepts to IVNs

- Commonalities
  - mostly-static topology
  - small buffers in network & destination
- Differences
  - traffic patterns
  - non-blocking fabric

Network-Scale Virtual Output Queues

- Applied to network endpoints
  - equivalent to chassis ingress
- One output queue per connection
  - “connection” = “two communicating software processes”
- Shaped output queues
- Shapers tuned to connection requirements

Source Endpoint Characteristics

- Unbounded source buffer size
- Unlimited enqueue bandwidth
- Multiple packets per queue entry

Separate Transmit Ports

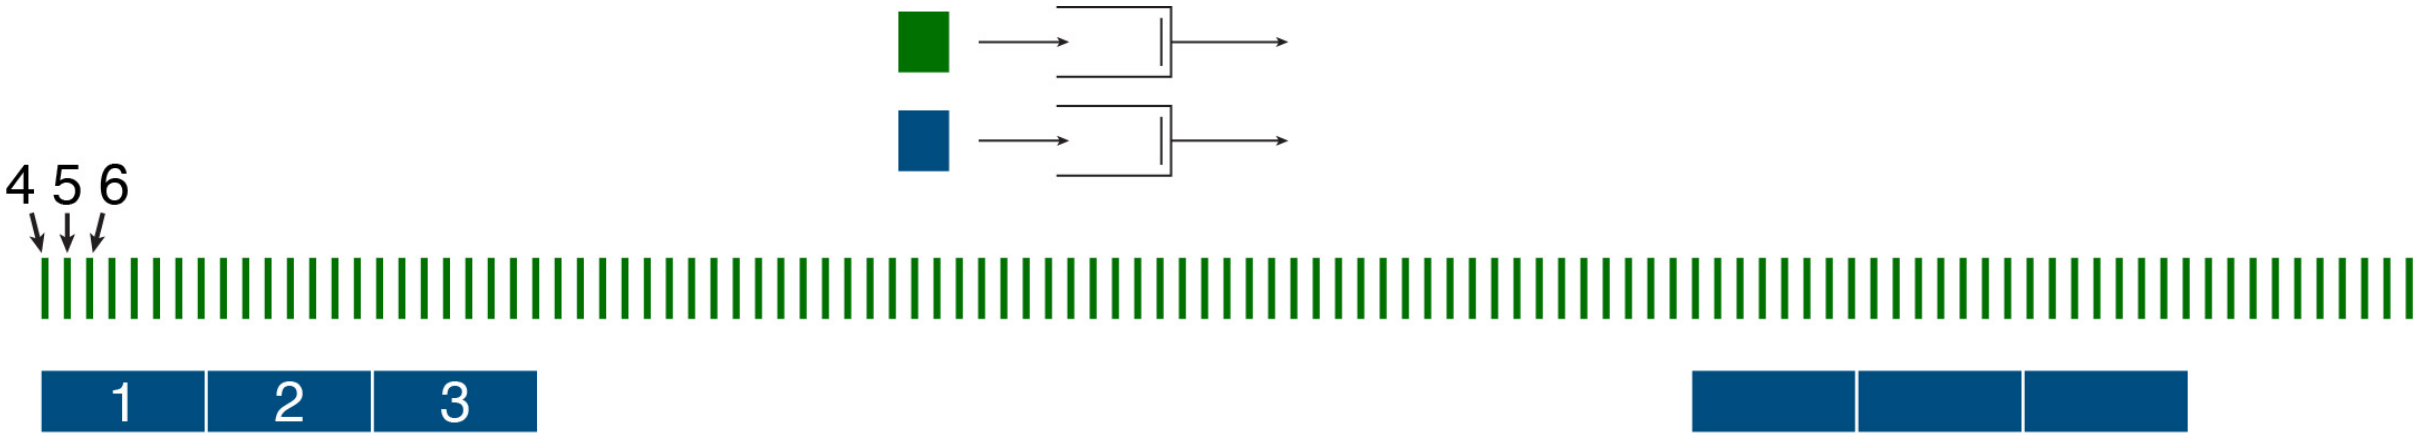

Shared Output Queue

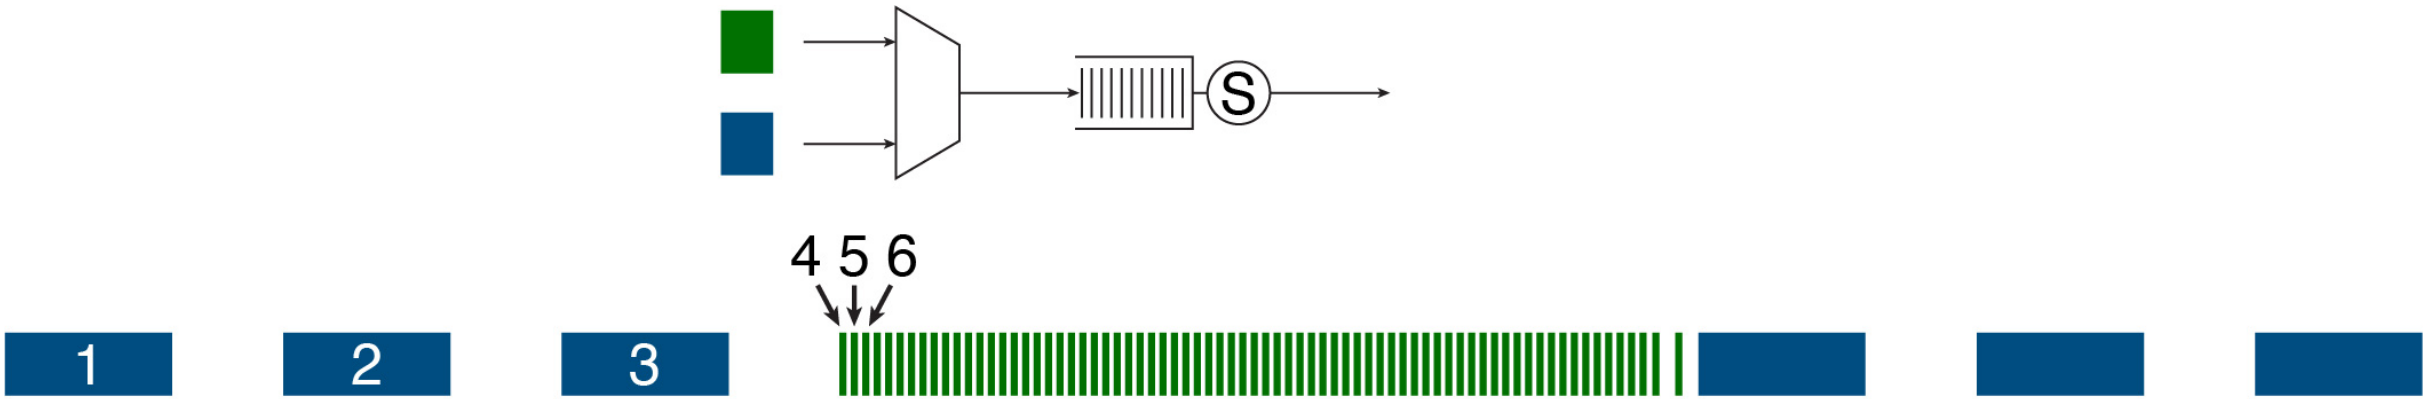

Separate Output Queues

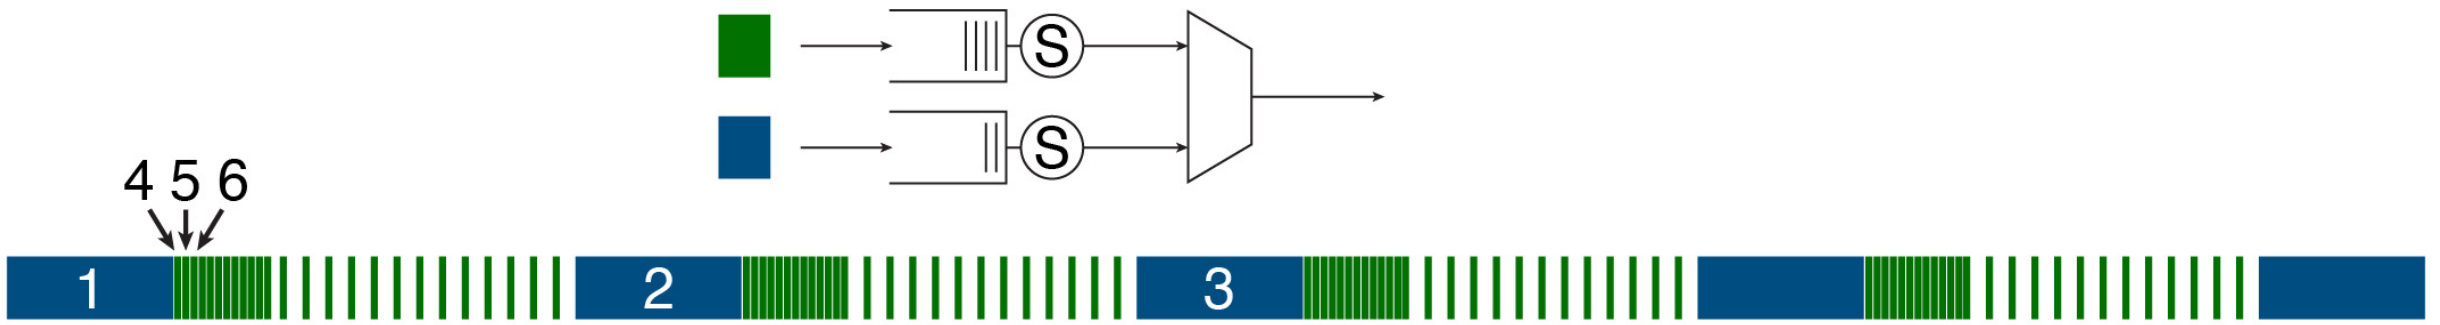

Comparison

shared output queue:

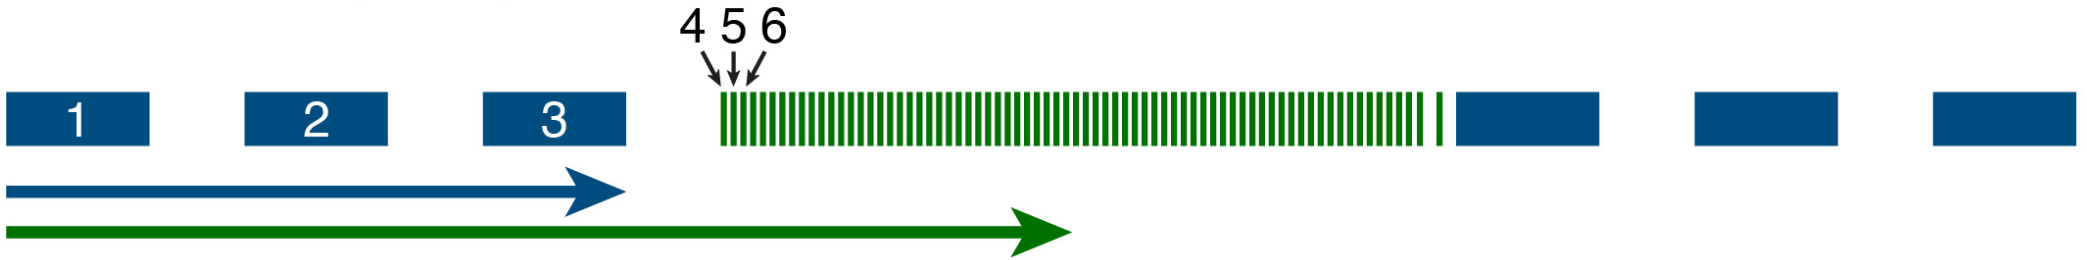

separate output queues:

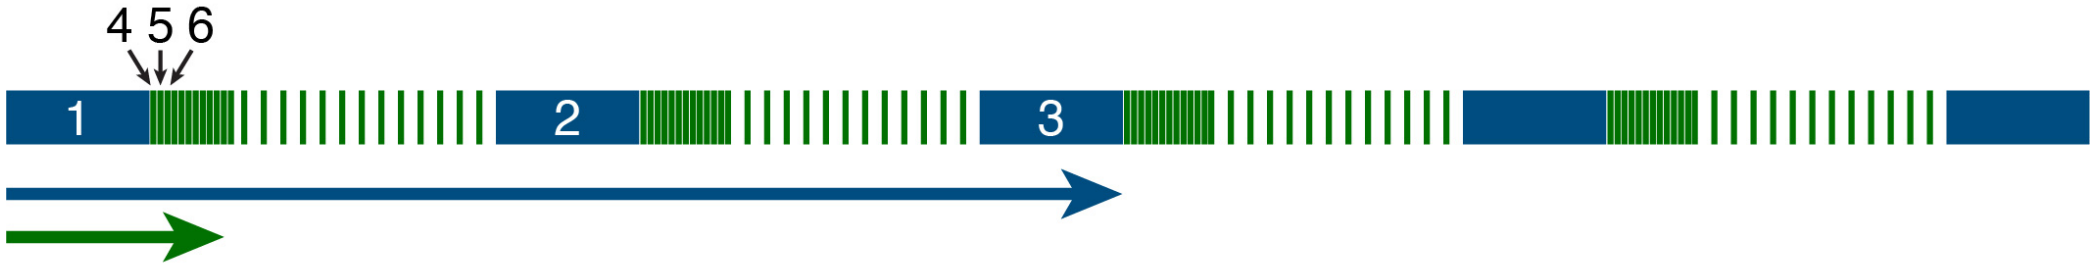

1.8x latency

0.2x latency

Downstream Bridge

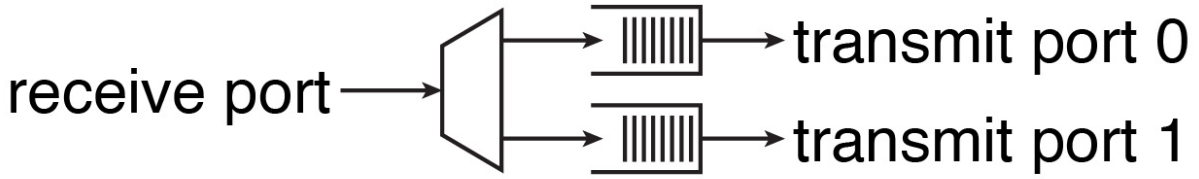

Ingress Endpoint Shared Output Queue:

Ingress Endpoint Separate Output Queues

Shaping Algorithms

- Asynchronous
  - ATS, CBS
  - doesn’t really matter
- Synchronous (TAS) adds unnecessary complexity
- Transmit rate ≤ configured rate
  - no queue depth growth
- Benefits greatly from network-wide bandwidth planning

VOQs and 802.1Q

- 8-queue limit
  - Tied to VLAN priority field
- VOQs operate at endpoints, not bridges
- 802.1Q does not apply

Conclusions

- Chassis and IVN share some characteristics
- Precision ingress rate controls prevent congestion
- VOQs solve head-of-line blocking
- VOQs reduce average latency

Questions?

Thank You