Design and Architecture of Data Center Networking Platforms

Lincoln Dale <ltd@cisco.com>
TME, Data Center Switching

Agenda

- Part I—Design Considerations
- Part II—Architectures
- Part III—Real World Implementation
- Part IV—The Future

Part I – Design Considerations

Part I—Design Considerations

- **Many Factors to Weigh**
- Buffering
- Multiprotocol
- One Plus One Is Not Necessarily Equal to Two
- Future (Investment Protection)

Many Factors to Weigh

Applicable to Any Switch/Router Design

- Standards requirements
- Market requirements
- Designability
- Silicon technology
- Processor technology
- Product is chassis **and** network
- Manufacturability
- Time to market
- Flexibility
- Budget

Many Factors to Weigh

Baseline Data Center Switch requirements

Data Plane

- Buffering
- No packet drop
- Throughput
- Port count
- Modular
- No single point of failure
- In-order delivery
- Future protocol compatibility

Control plane

- Modular
- Restartable (including active-active state handling)
- Non-disruptive code load & activation
- No single point of failure
- Scaleable
- Unit Testable
- Future protocol compatibility

Part I—Design Considerations

- Many Factors to Weigh
- **Buffering**
- Multiprotocol
- One Plus One Is Not Necessarily Equal to Two
- Future (Investment Protection)

Buffering

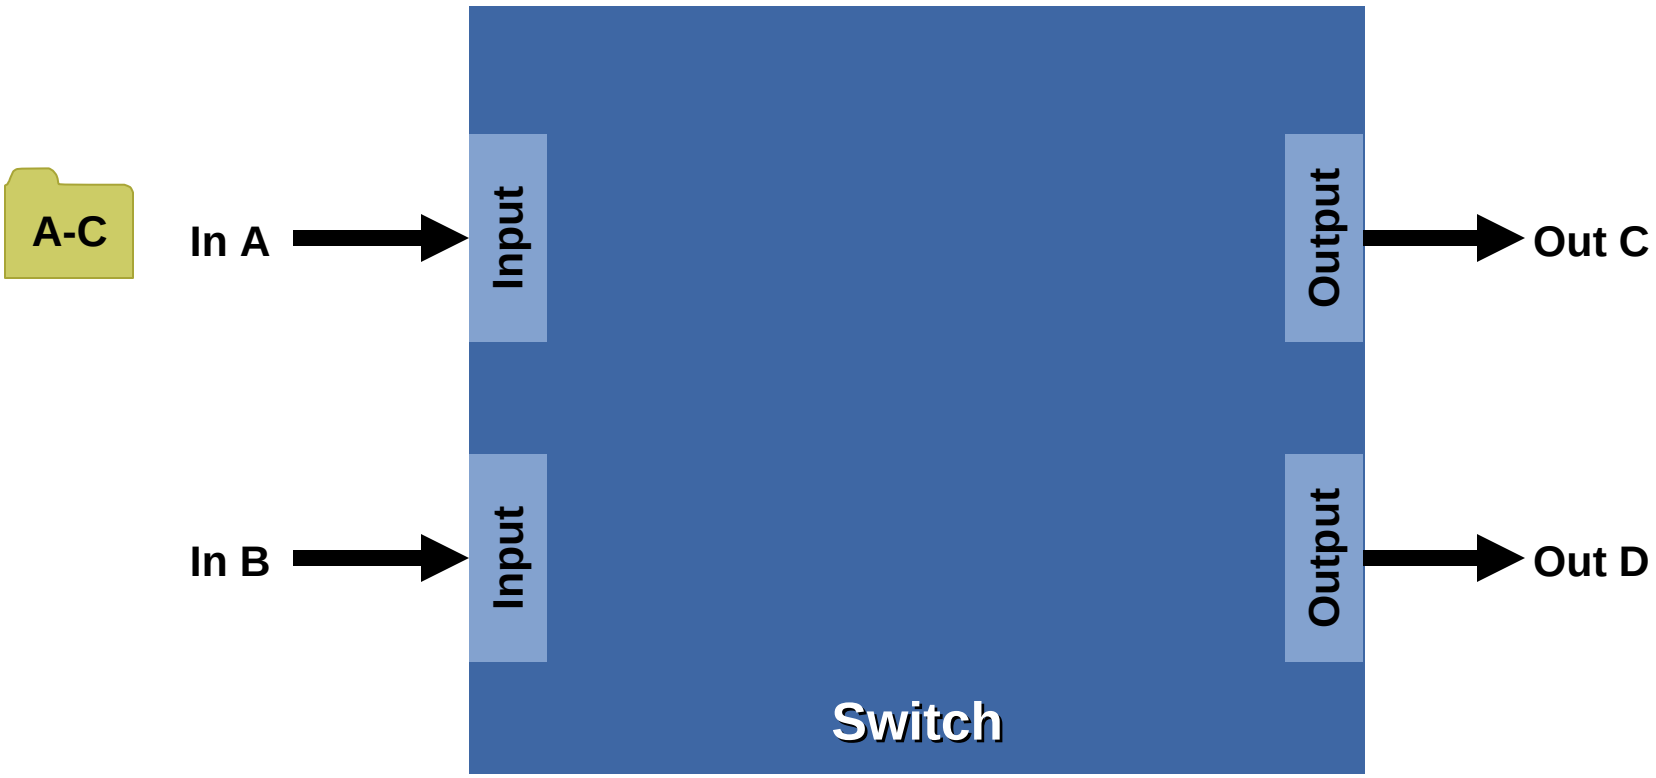

Buffering

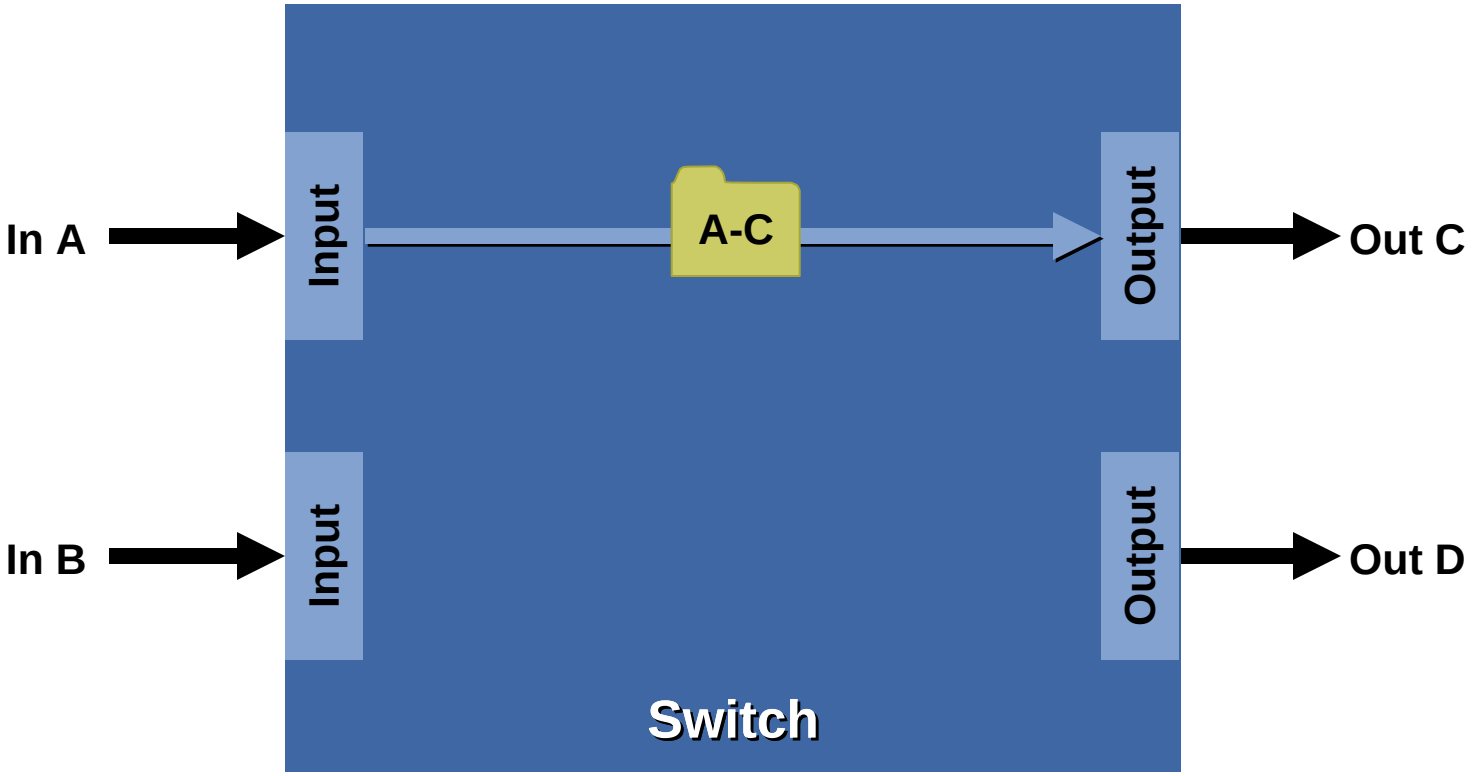

Buffering

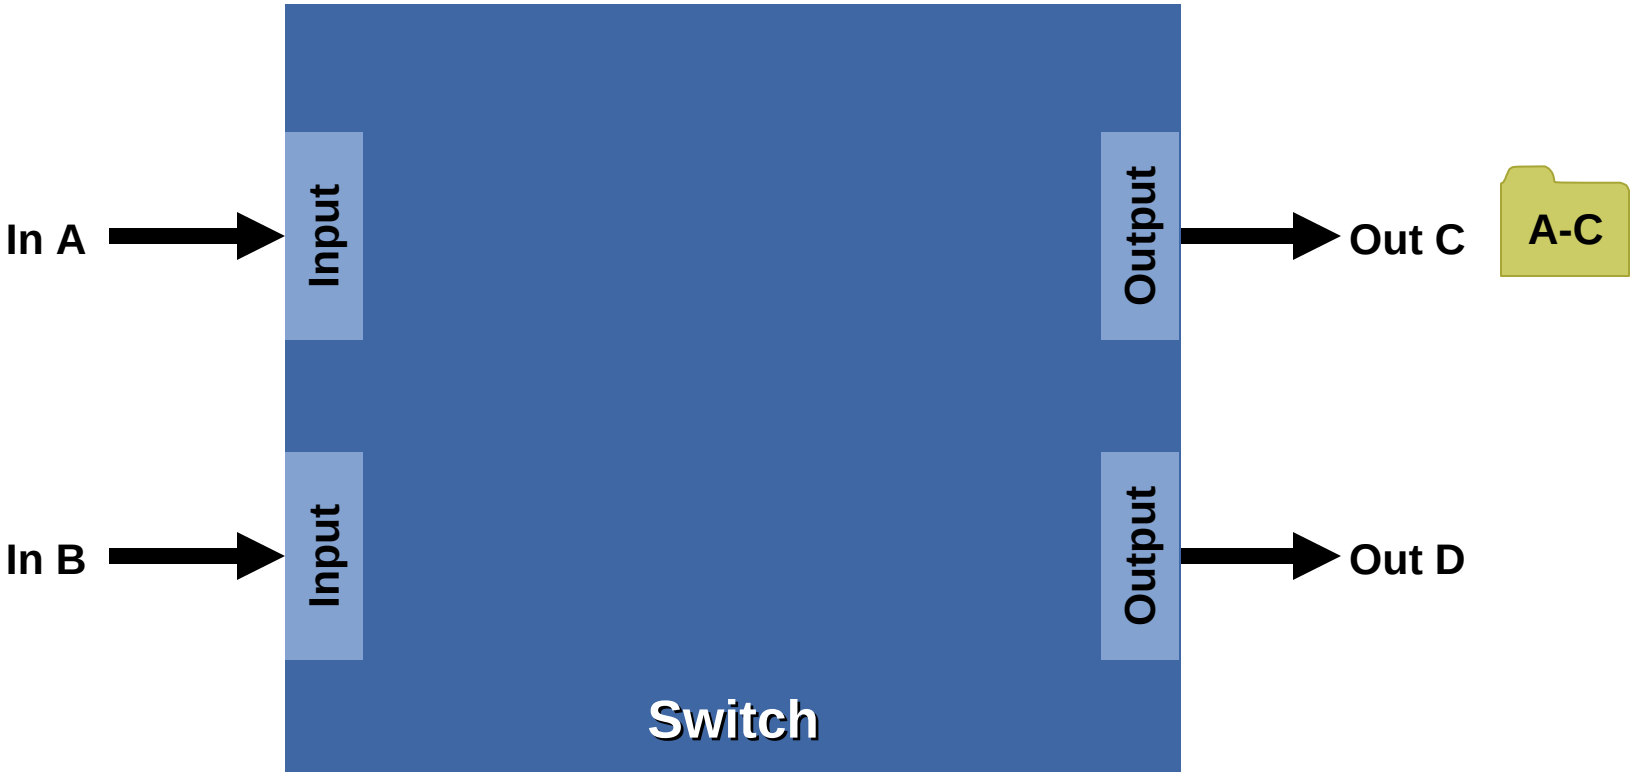

Buffering

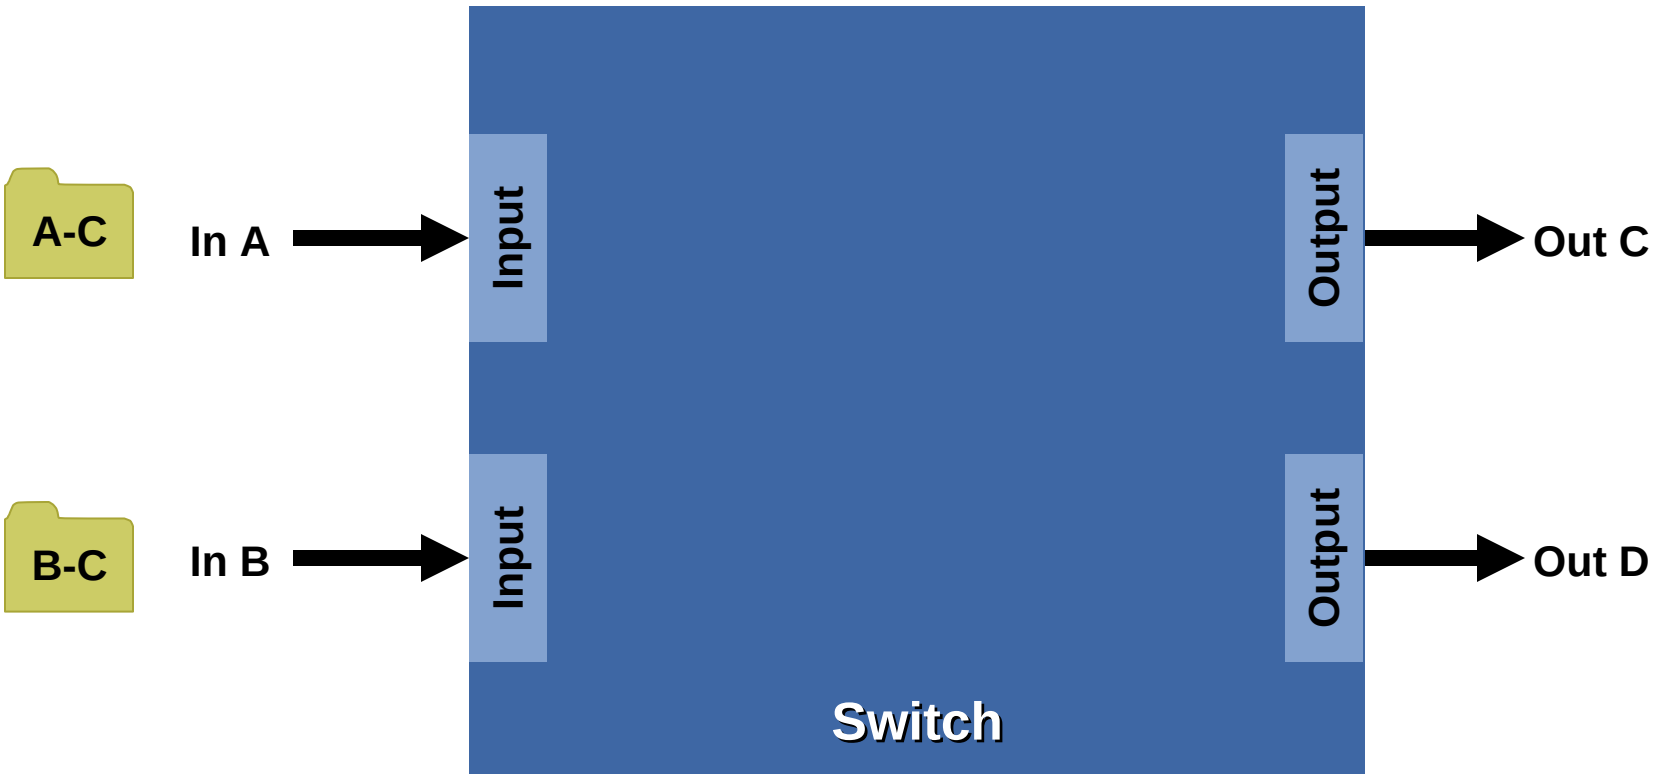

Buffering

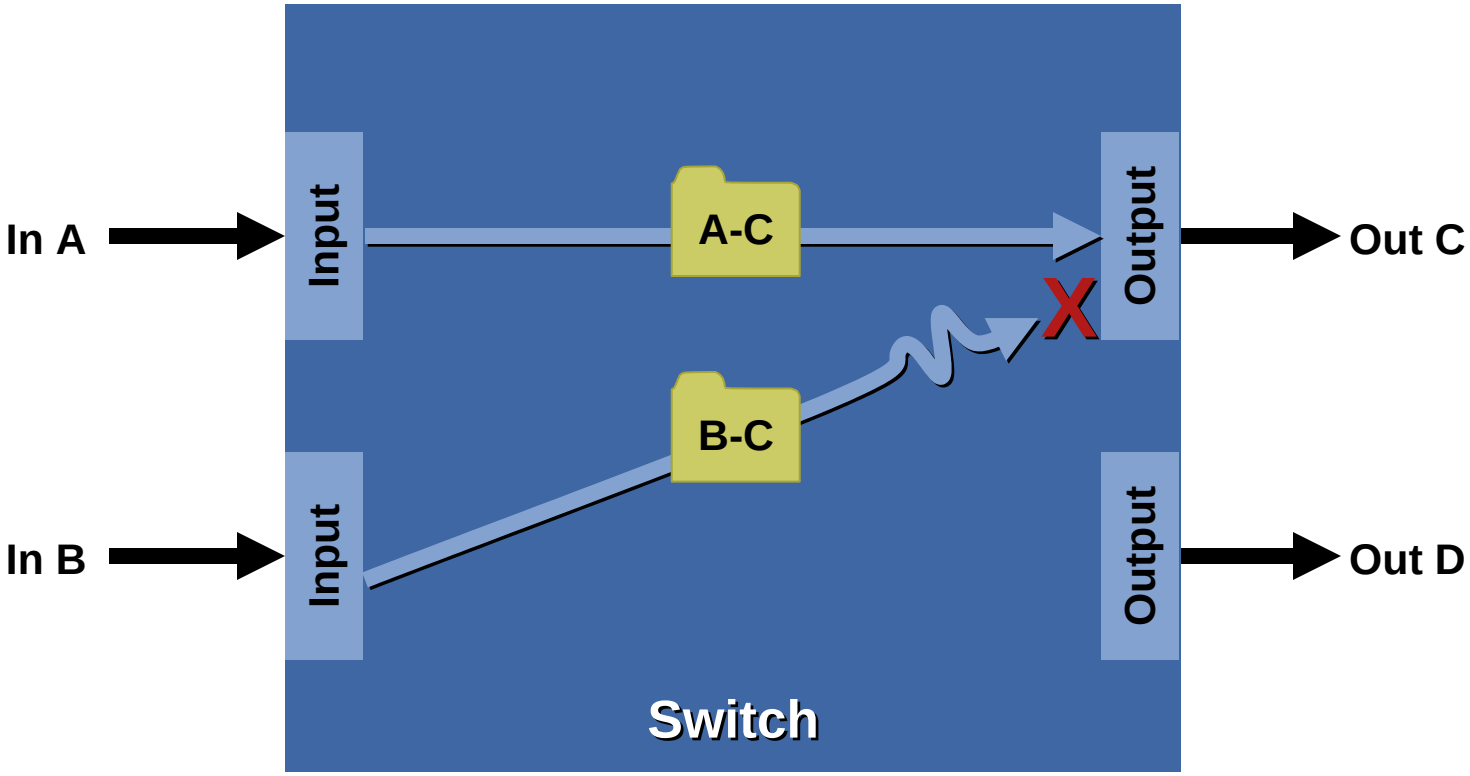

Buffering

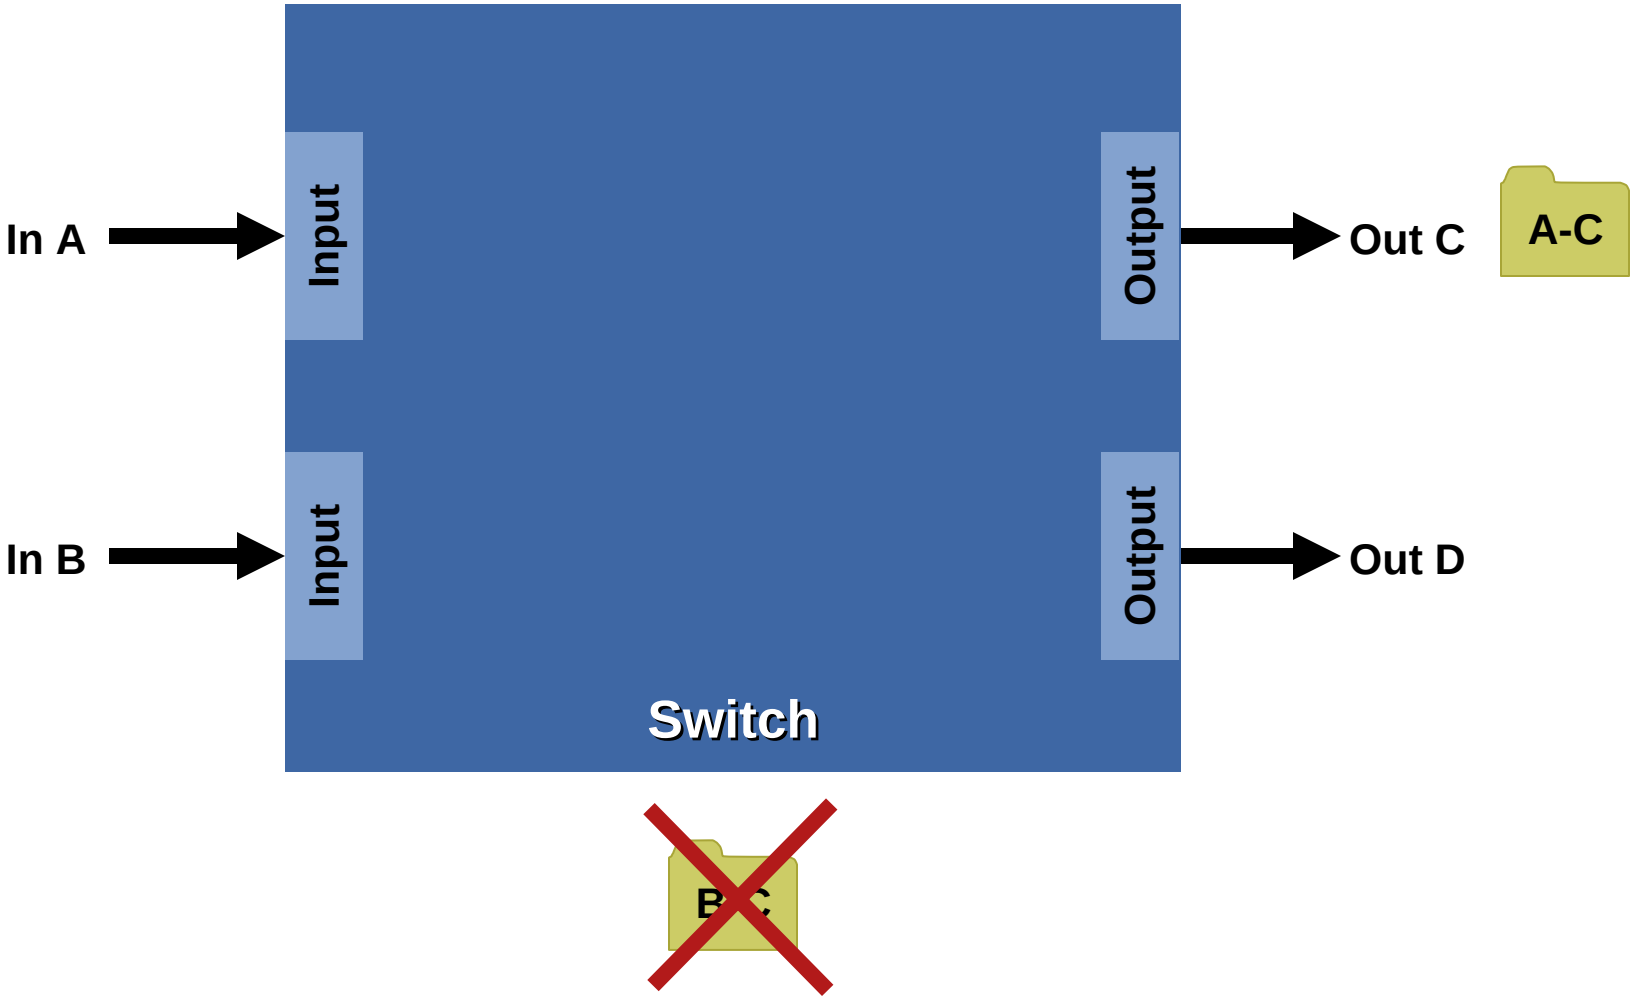

Buffering

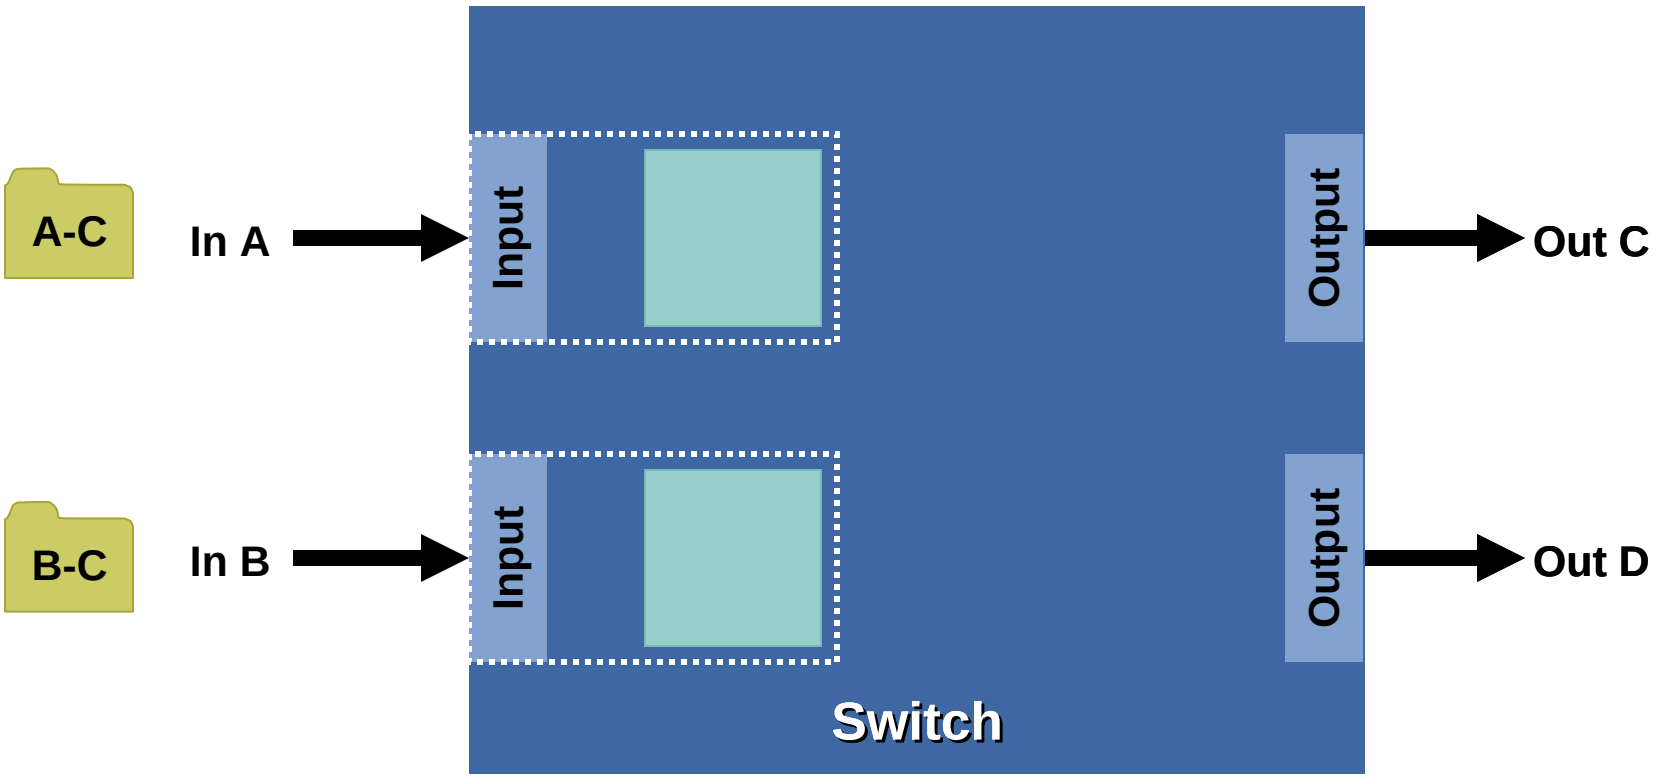

Buffering

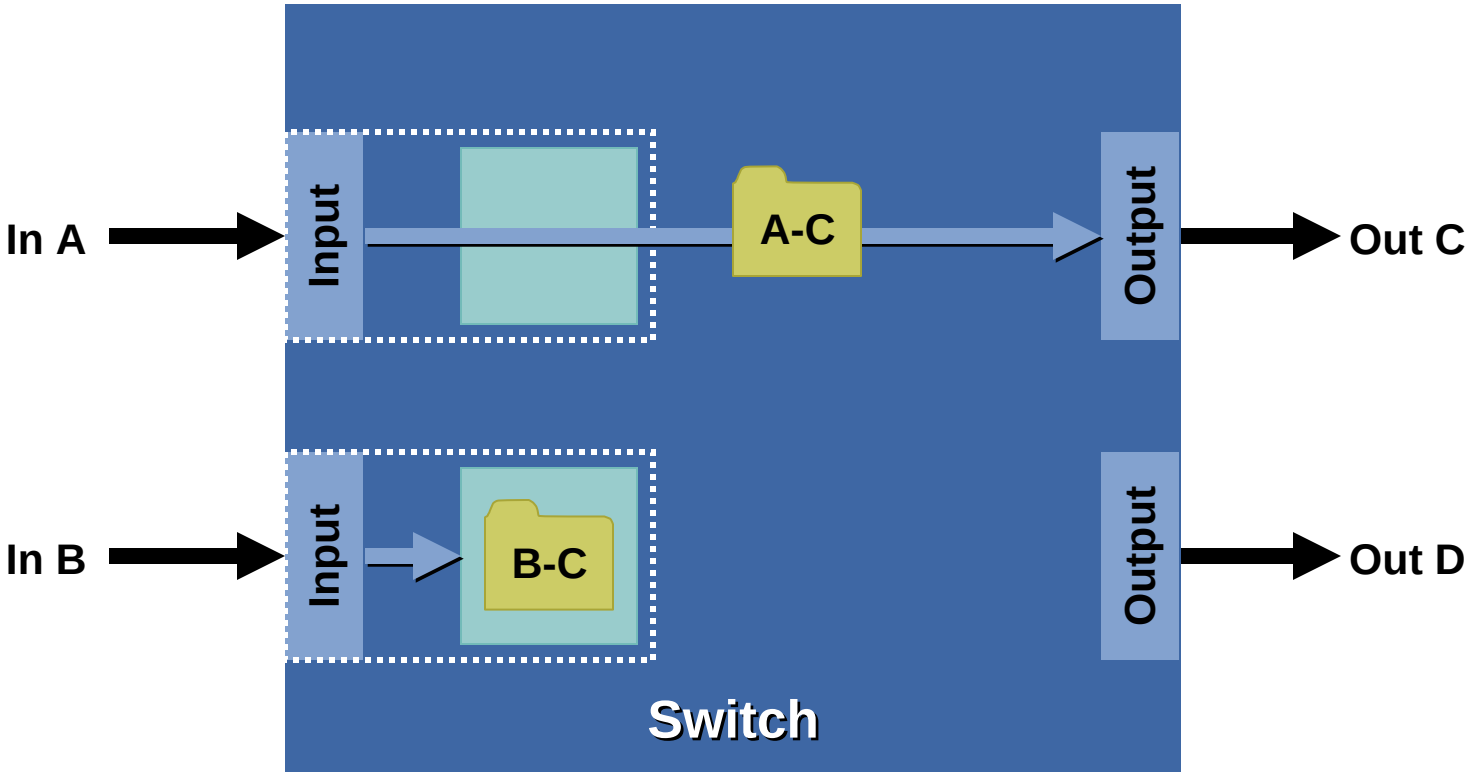

Buffering

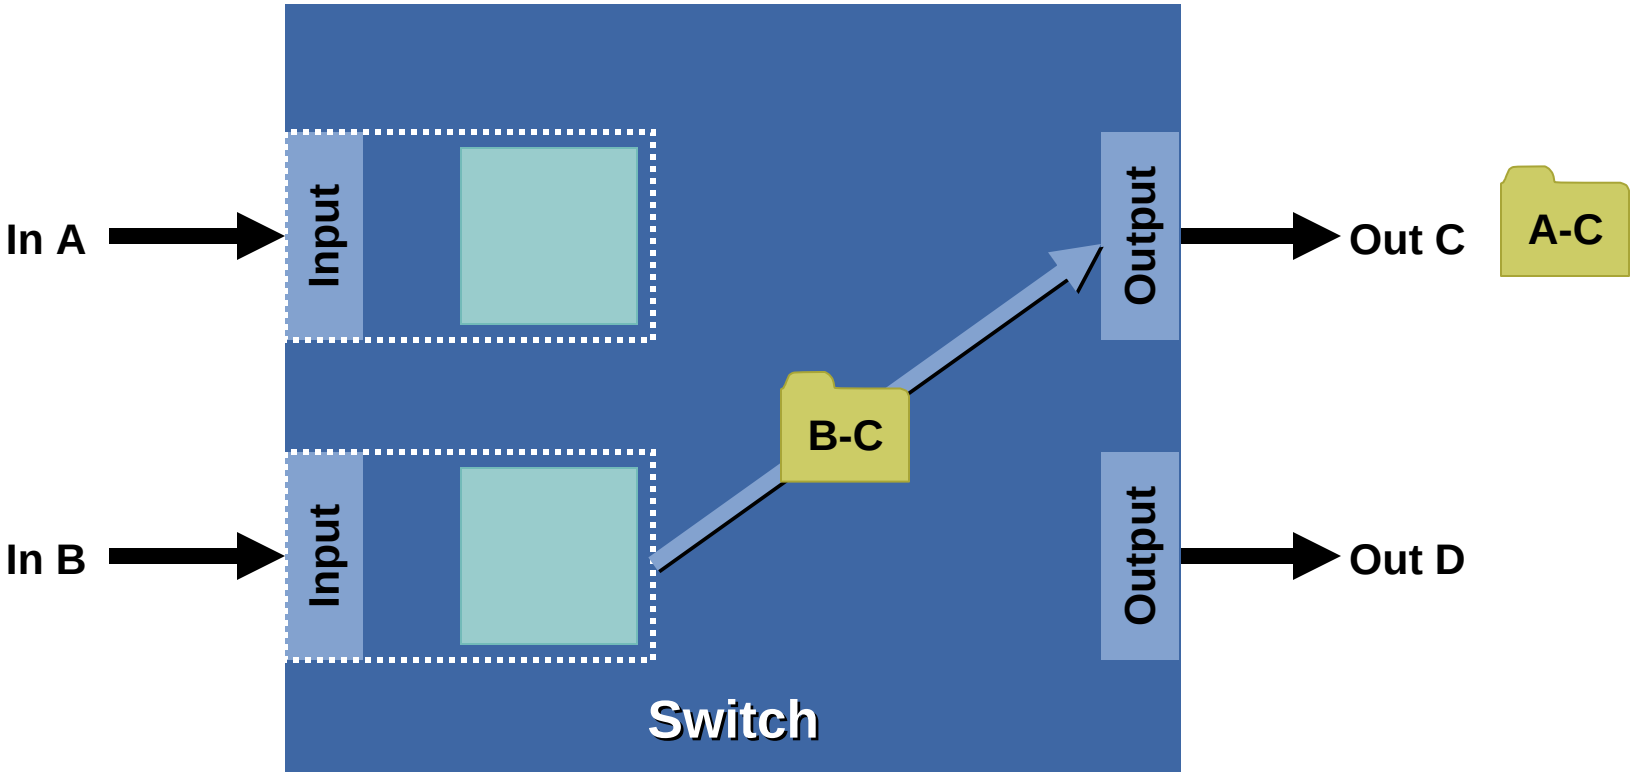

Buffering

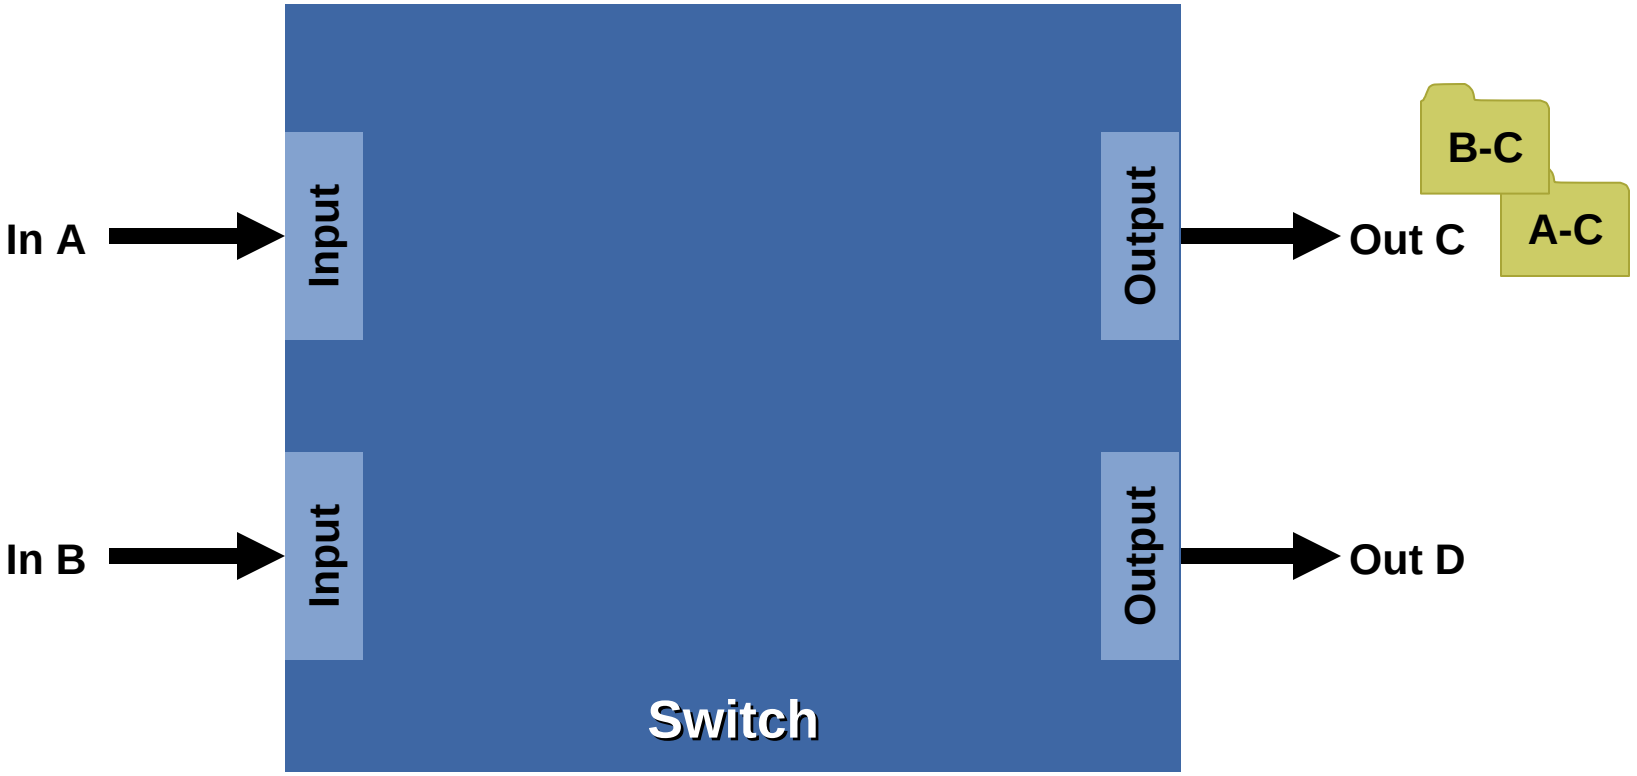

Part I—Design Considerations

- Many Factors to Weigh
- Buffering
- **Multiprotocol**
- One Plus One Is Not Necessarily Equal to Two
- Future (Investment Protection)

Multi Protocol

- **Protocol Agnostic**
- **Protocol Translation**
- **Drop Behavior**
- **Shared Uplinks**
- **Multiple Topologies**

- 10Base{2,5,T} / 100Base{S,T,TX} / 1000BaseT
- 10GbE
- Fibre Channel
- FICON
- ESCON
- Token Ring
- 100BaseVG

- IP & TCP
- IPv4/IPv6
- 802.1q VLANs
- MPLS / MPLS {P,PE} / MPLE {CE,..}, ..
- VPLS
- q-in-q VLAN encapsulation
- 802.2 LLC/SNAP
- ATM
- Frame Relay
- DOCSIS
- AToM

- IP/Ethernet
- 802.11 “WiFi”
- 802.16 “WiMax”
- FDDI
- Vines
- XNS
- AppleTalk
- IPX/SPX
- Netbeui
- DECnet
- OSI/ISO
- ArcNet
- L2TPv3

Part I—Design Considerations

- Many Factors to Weigh
- Buffering
- Multiprotocol
- **One Plus One Is Not Necessarily Equal to Two**
- Future (Investment Protection)

One Plus One Does Not Equal Two

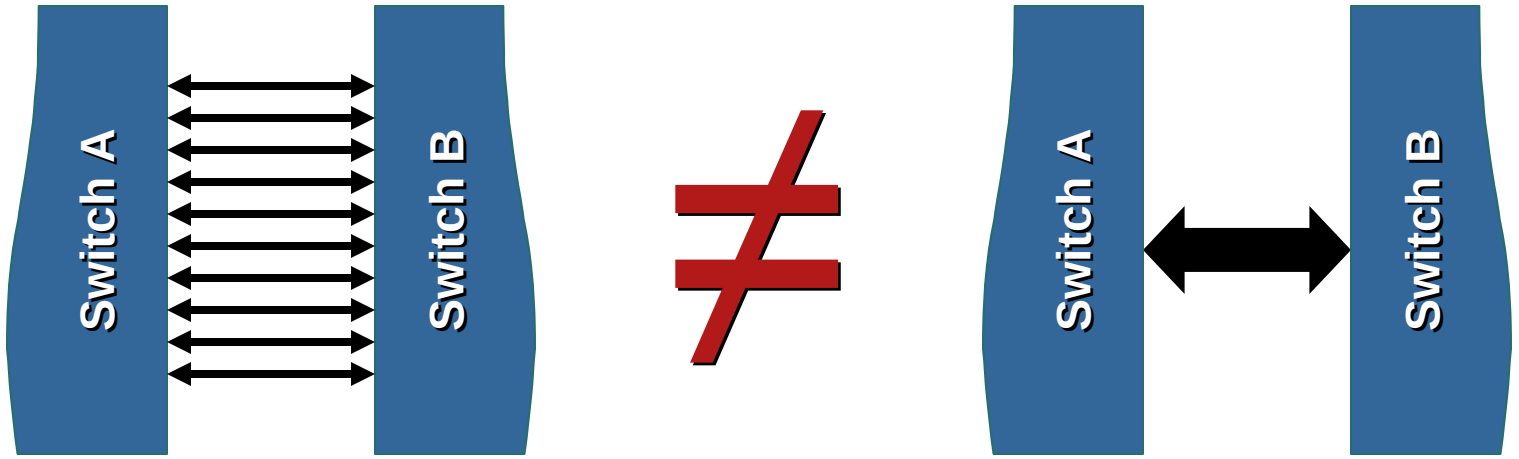

10 Links @ 1Gbps Each
Bandwidth = 10Gbps
Flow Bandwidth = 1Gbps
Serialization Delay = 20uS

1 Link @ 10Gbps Each
Bandwidth = 10Gbps
Flow Bandwidth = 10Gbps
Serialization Delay = 2uS

Part II – Architectures

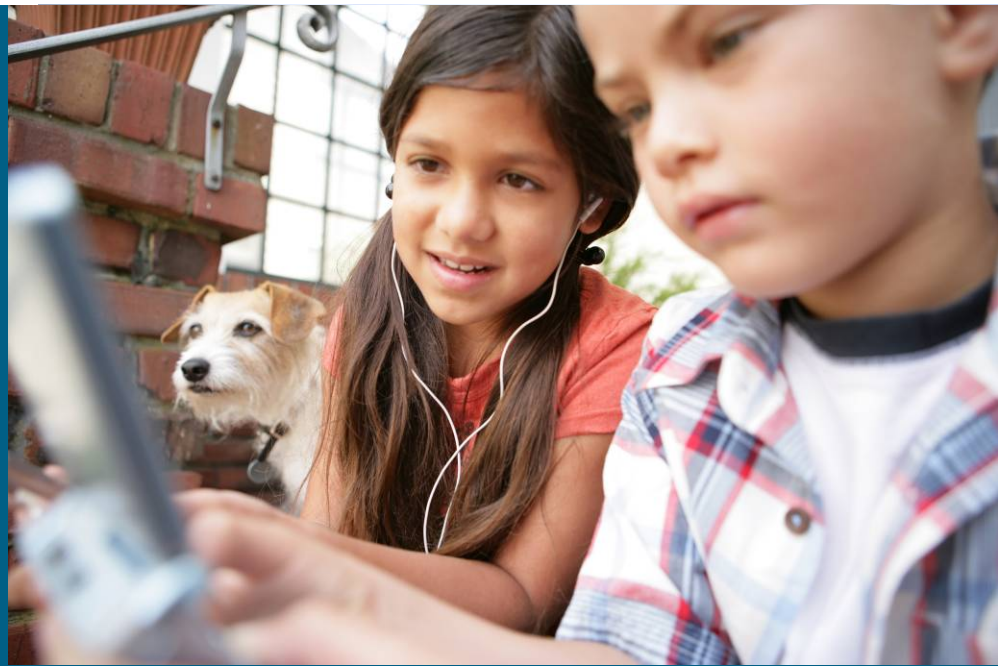

Part II—Architectures

- Centralized Memory
  - Switch on a Chip / “System on Chip” (SoC)
- Bus
- Mesh
  - Centralized memory in a mesh
- Two-Tier / Banyan / Clos
  - Centralized memory in two tiers
- Cross-Bar

Part II—Architectures

- **Centralized Memory**
  - **Switch on a chip / “System on Chip” (SoC)**
- Bus
- Mesh
  - Centralized memory in a mesh
- Two-Tier / Banyan / Clos
  - Centralized memory in two tiers
- Cross-Bar

Centralized Memory

- Scalability limited by memory bandwidth/size
- Typically optimized for fixed configuration—inflexible
- Port count 8–32
- Cost effective with small port counts
- Often used as building block
- Almost required for Blade Server Applications

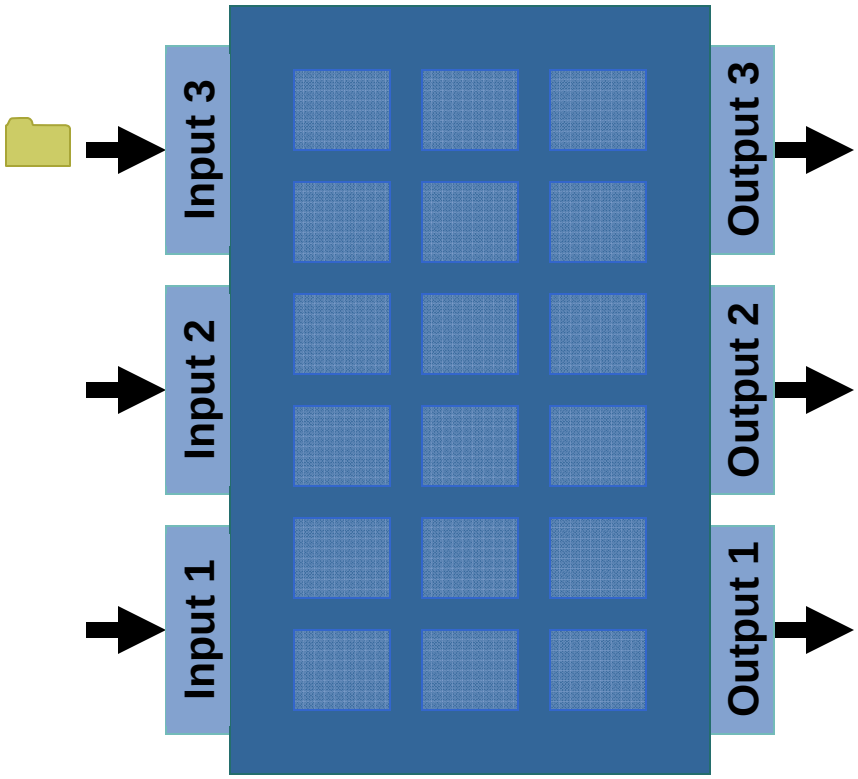

Centralized Memory

- Scalability limited by memory bandwidth/size
- Typically optimized for fixed configuration— inflexible
- Port count 8–32
- Cost effective with small port counts
- Often used as building block
- Almost required for Blade Server Applications

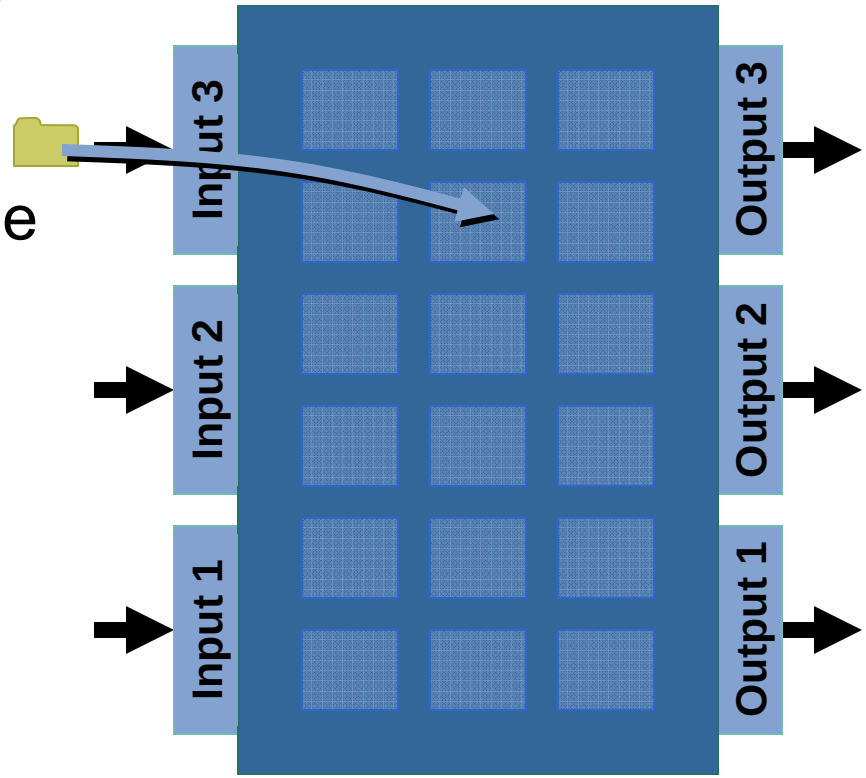

Centralized Memory

- Scalability limited by memory bandwidth/size
- Typically optimized for fixed configuration— inflexible
- Port count 8–32
- Cost effective with small port counts
- Often used as building block
- Almost required for Blade Server Applications

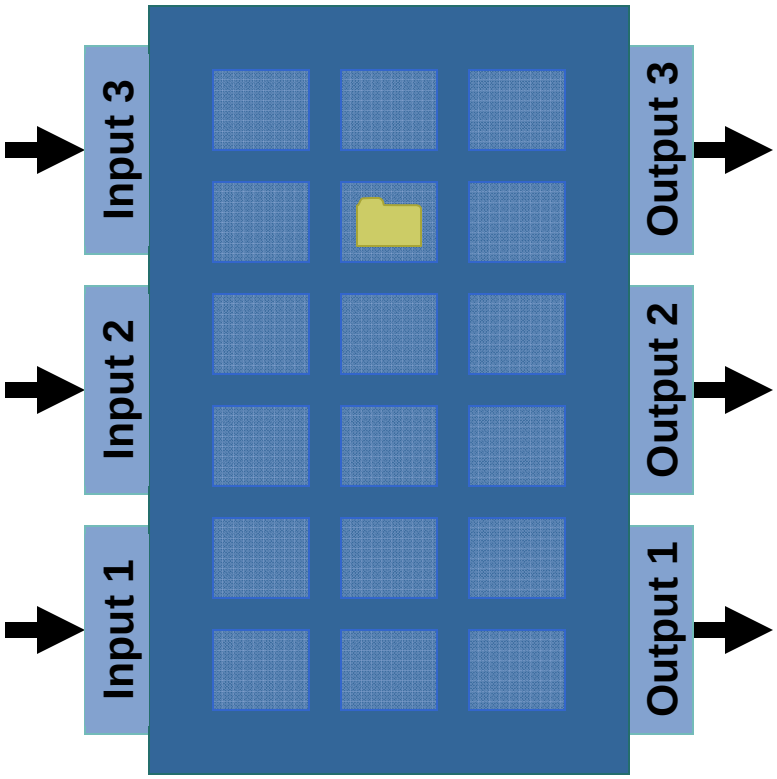

Centralized Memory

- Scalability limited by memory bandwidth/size
- Typically optimized for fixed configuration— inflexible
- Port count 8–32
- Cost effective with small port counts
- Often used as building block
- Almost required for Blade Server Applications

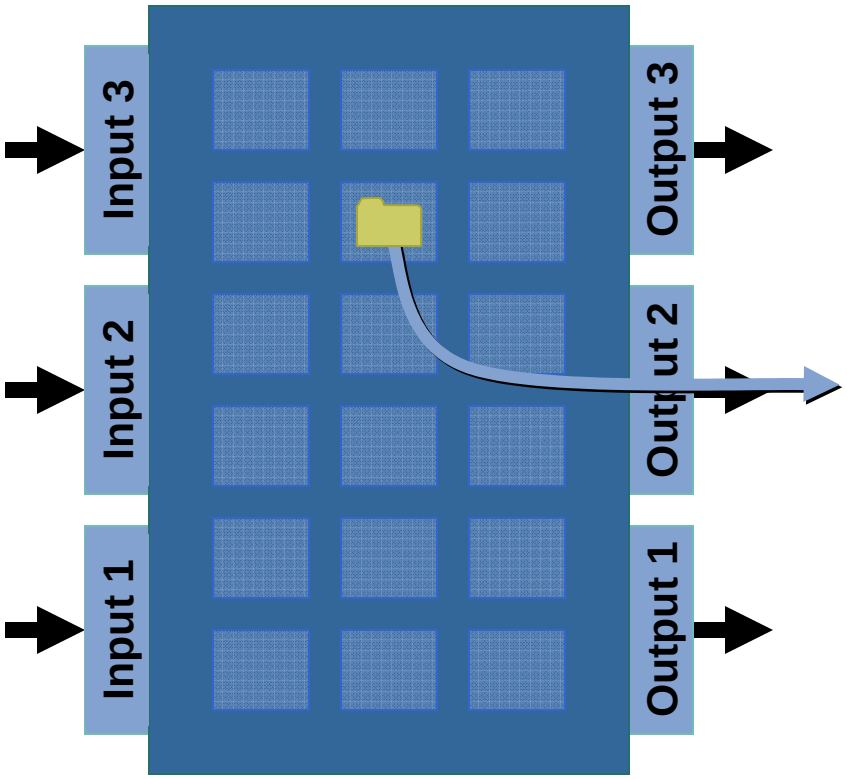

Centralized Memory

- Scalability limited by memory bandwidth/size
- Typically optimized for fixed configuration— inflexible
- Port count 8–32
- Cost effective with small port counts
- Often used as building block
- Almost required for Blade Server Applications

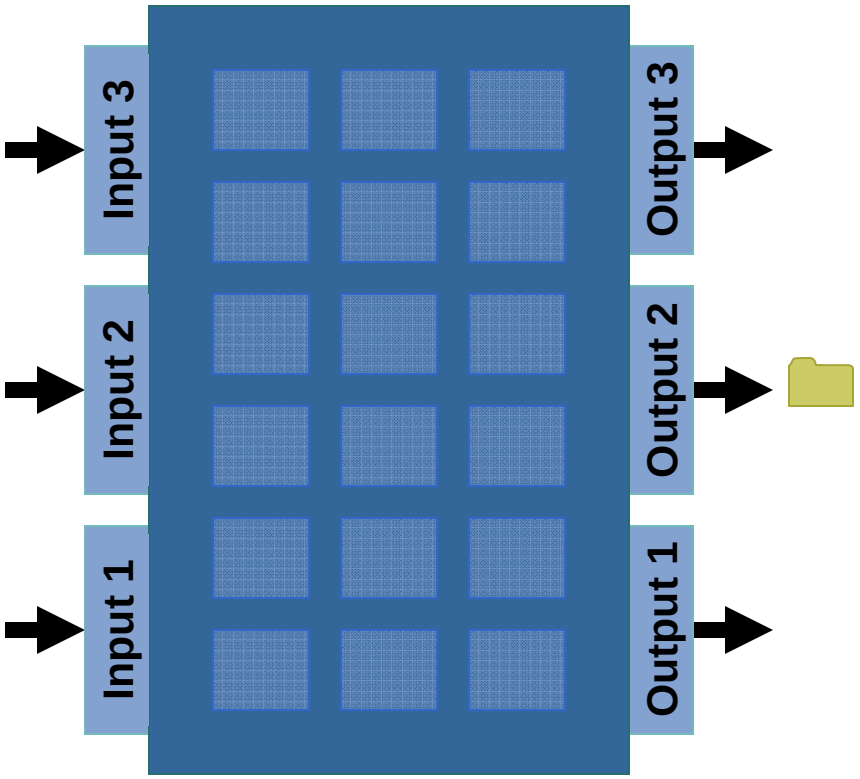

Centralized Memory

- Scalability limited by memory bandwidth/size
- Typically optimized for fixed configuration— inflexible
- Port count 8–32
- Cost effective with small port counts
- Often used as building block
- Almost required for Blade Server Applications

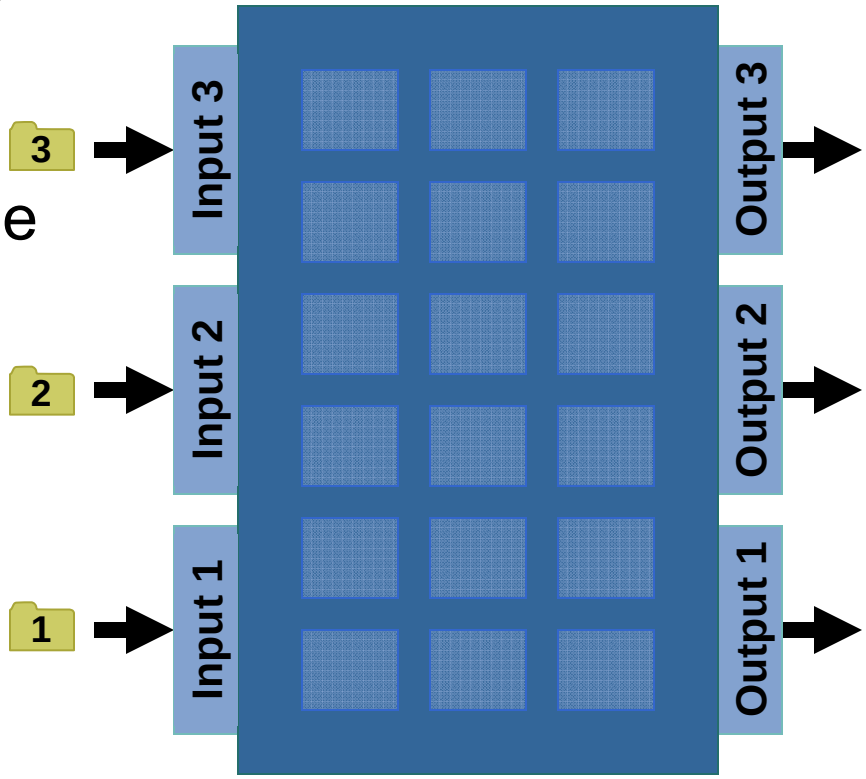

Centralized Memory

- Scalability limited by memory bandwidth/size
- Typically optimized for fixed configuration— inflexible
- Port count 8–32
- Cost effective with small port counts
- Often used as building block
- Almost required for Blade Server Applications

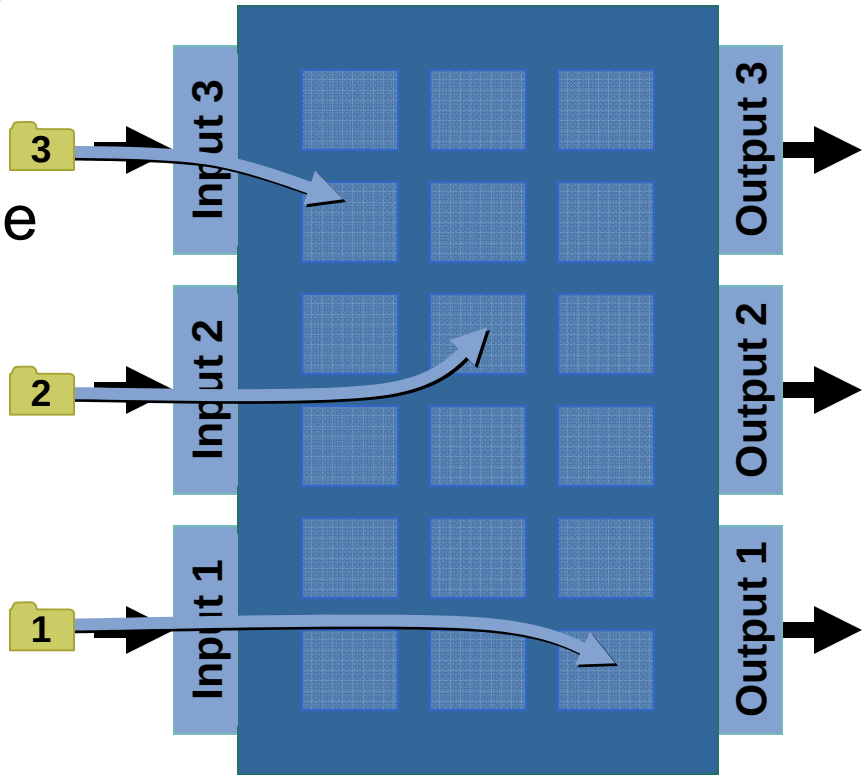

Centralized Memory

- Scalability limited by memory bandwidth/size
- Typically optimized for fixed configuration— inflexible
- Port count 8–32
- Cost effective with small port counts
- Often used as building block
- Almost required for Blade Server Applications

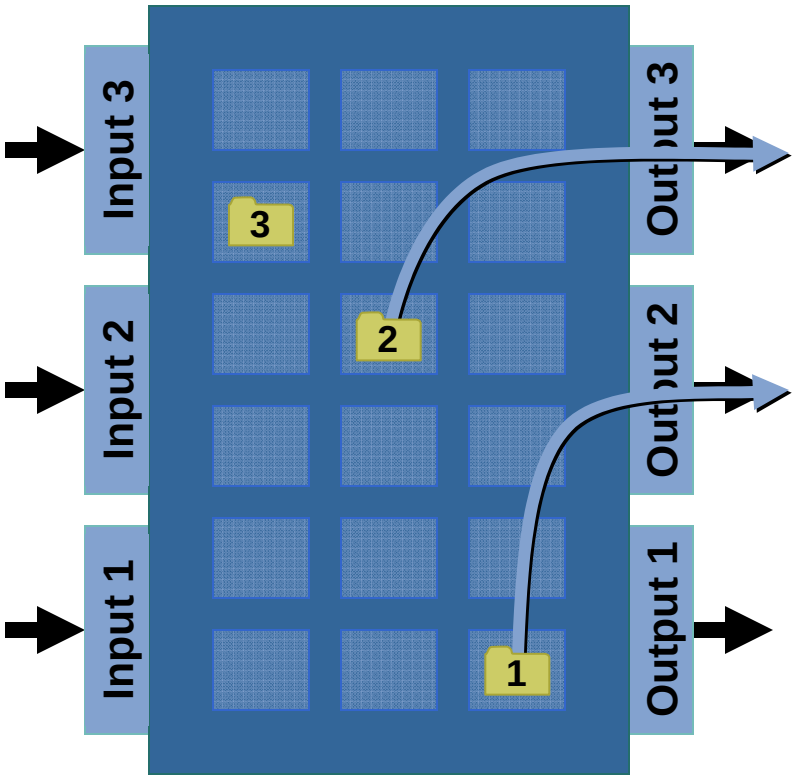

Centralized Memory

- Scalability limited by memory bandwidth/size
- Typically optimized for fixed configuration— inflexible
- Port count 8–32
- Cost effective with small port counts
- Often used as building block
- Almost required for Blade Server Applications

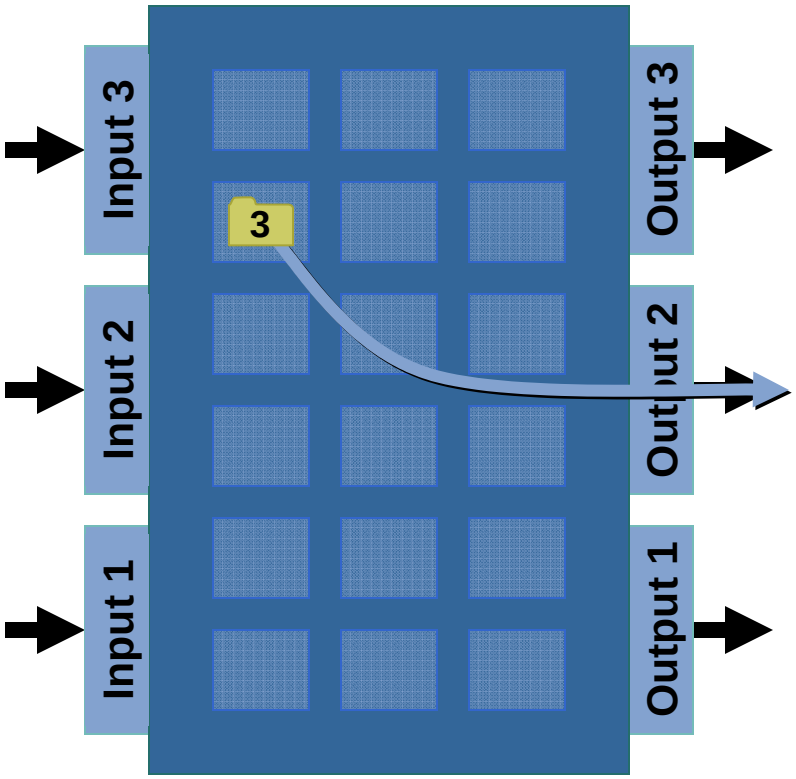

Centralized Memory

- Scalability limited by memory bandwidth/size
- Typically optimized for fixed configuration— inflexible
- Port count 8–32
- Cost effective with small port counts
- Often used as building block
- Almost required for Blade Server Applications

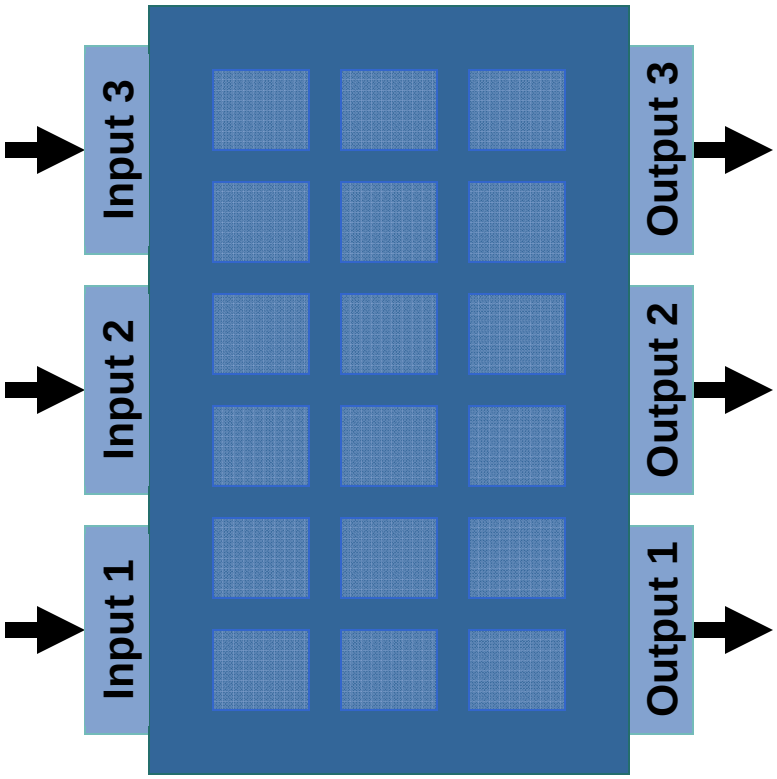

Part II—Architectures

- Centralized Memory
  - “Switch on a chip” / System on Chip (SoC)
- **Bus**
- Mesh
  - Centralized memory in a mesh
- Two-Tier / Banyan / Clos
  - Centralized memory in two tiers
- Cross-Bar

Bus

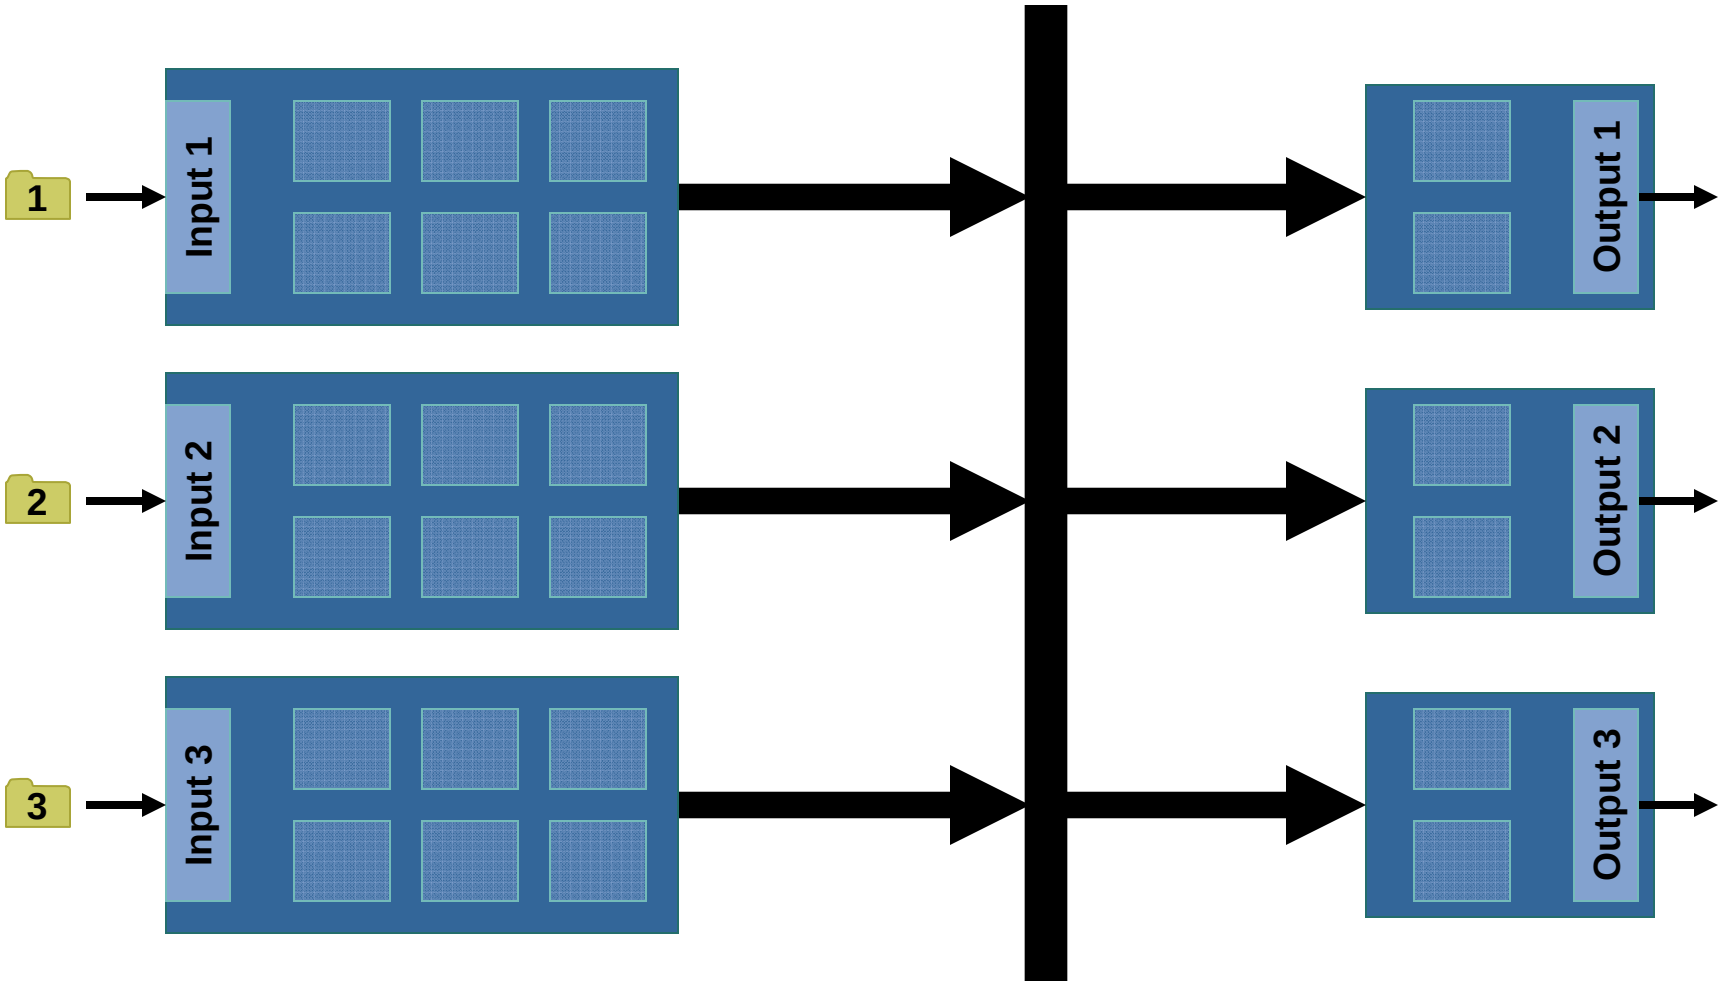

Bus

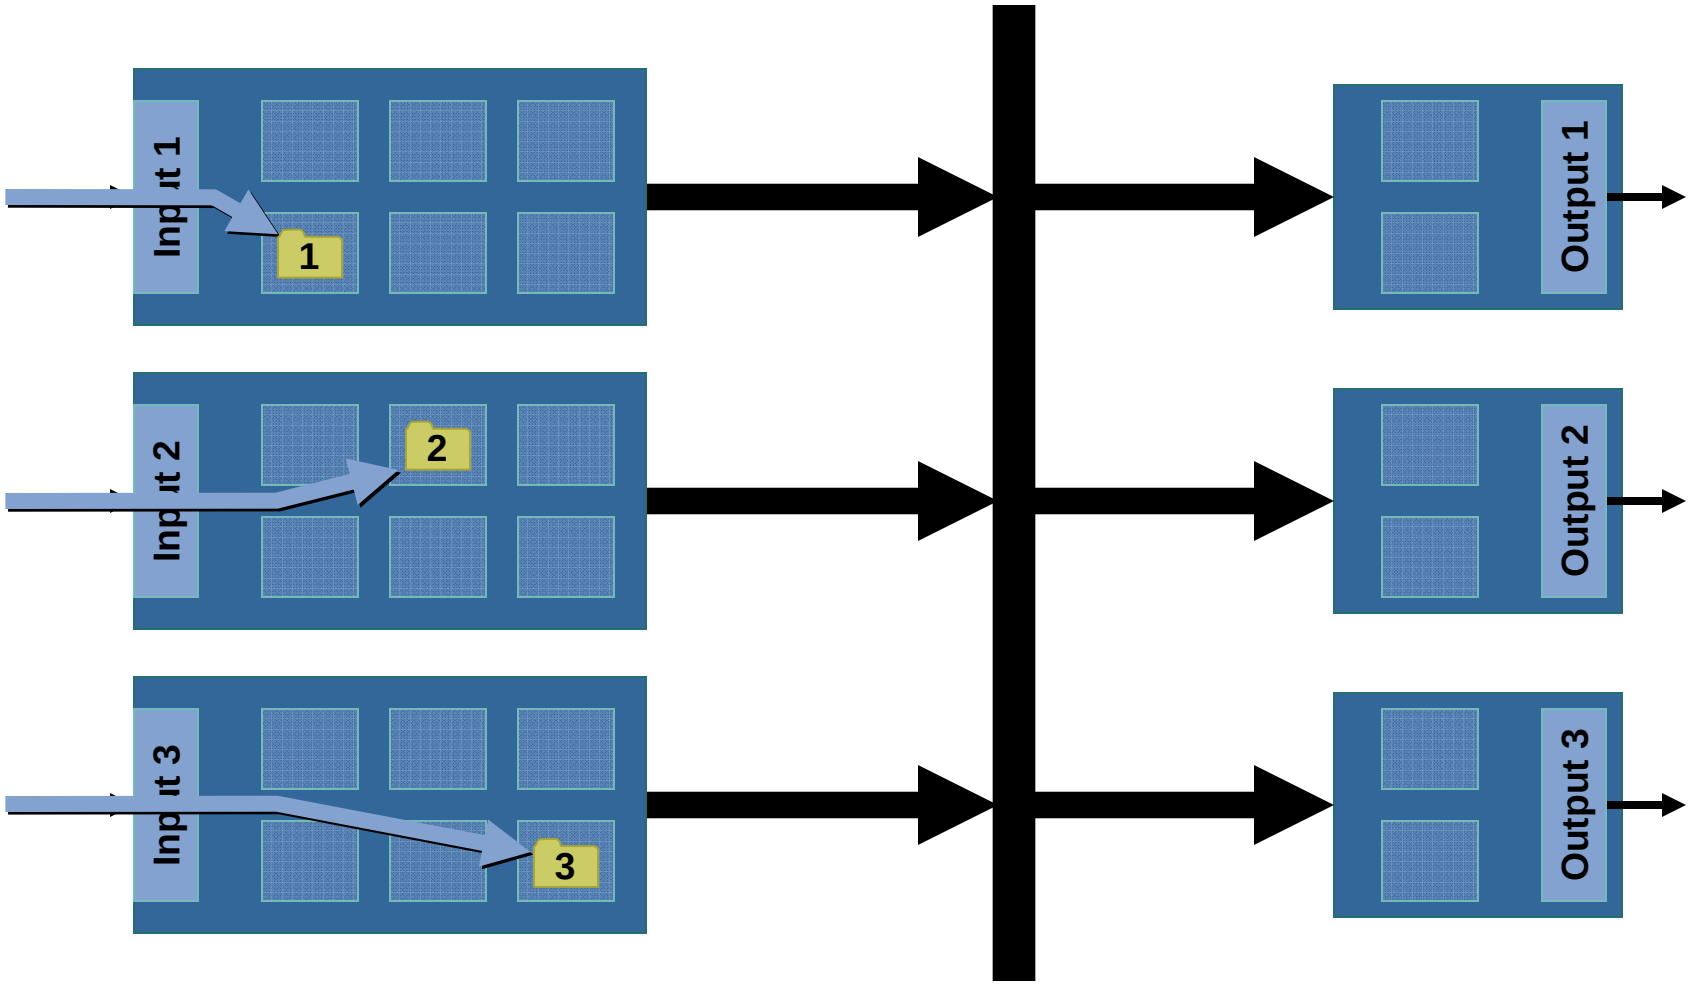

Bus

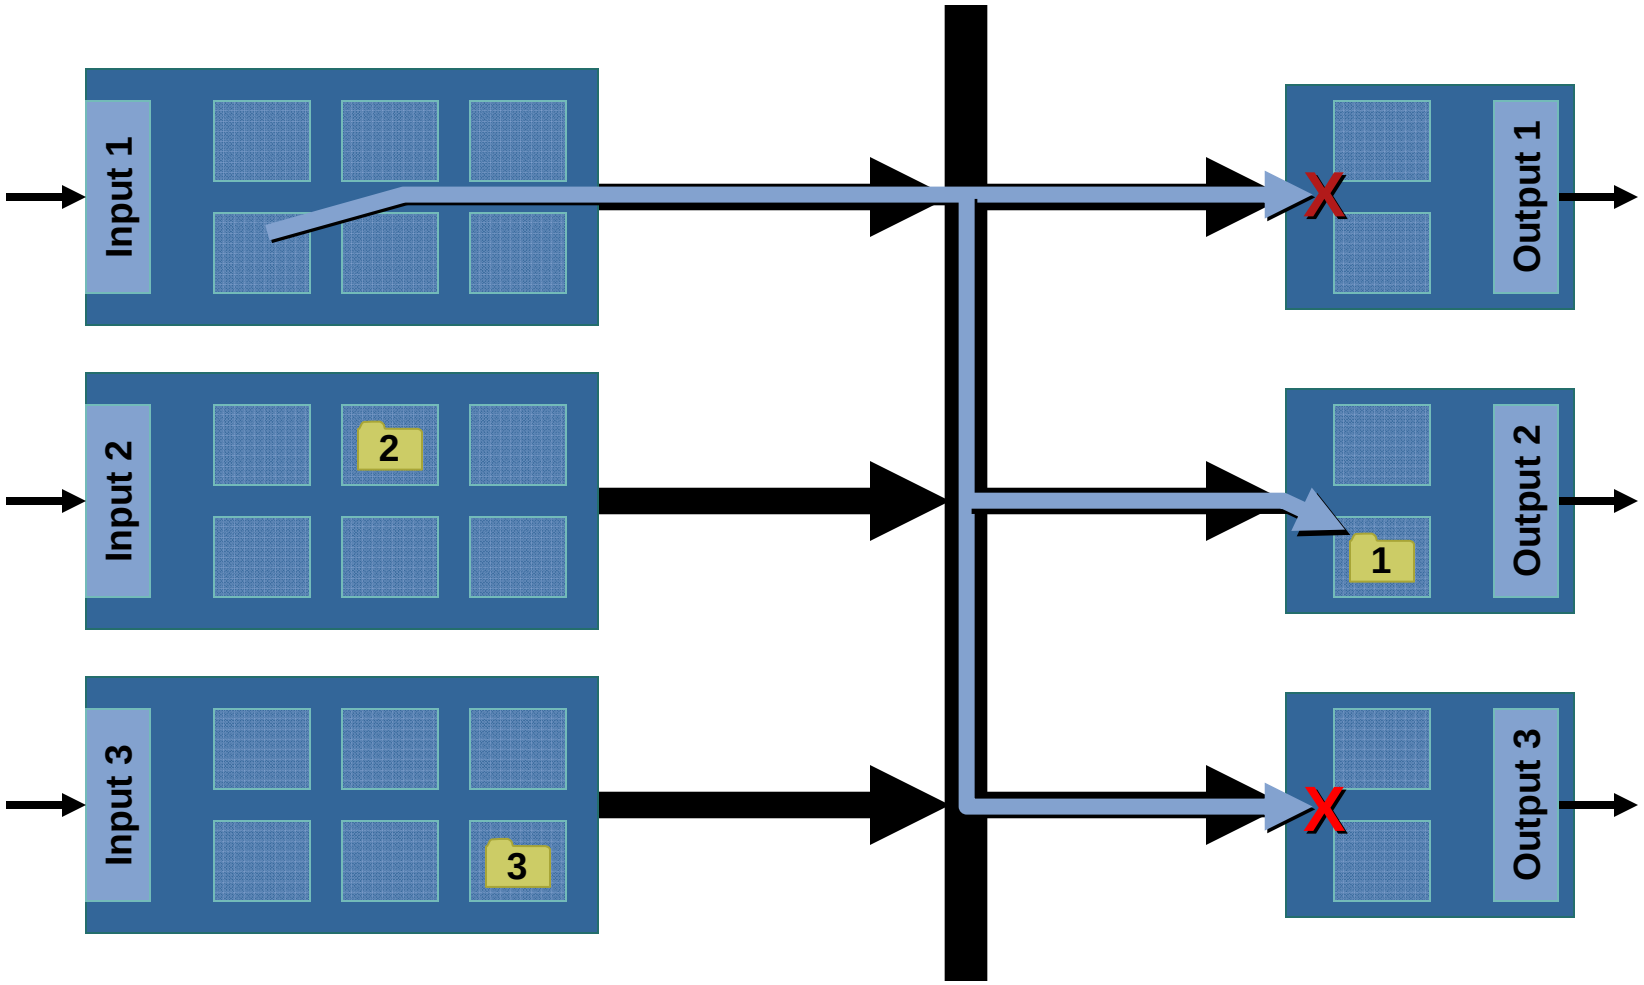

Bus

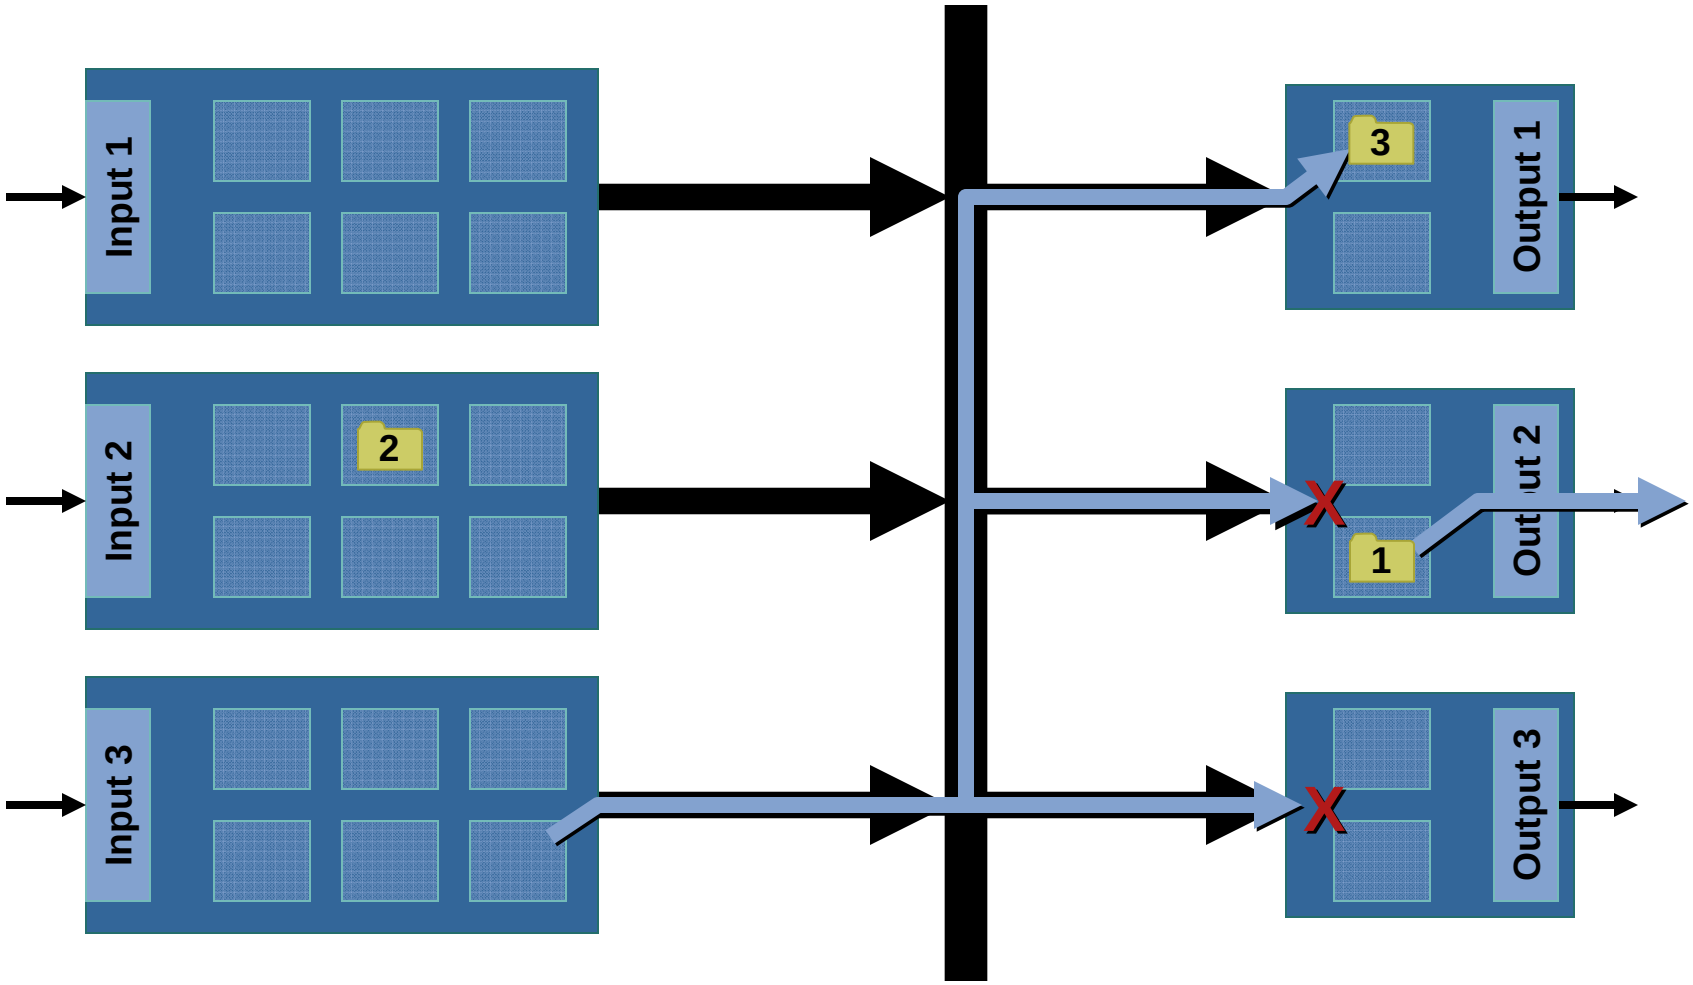

Bus

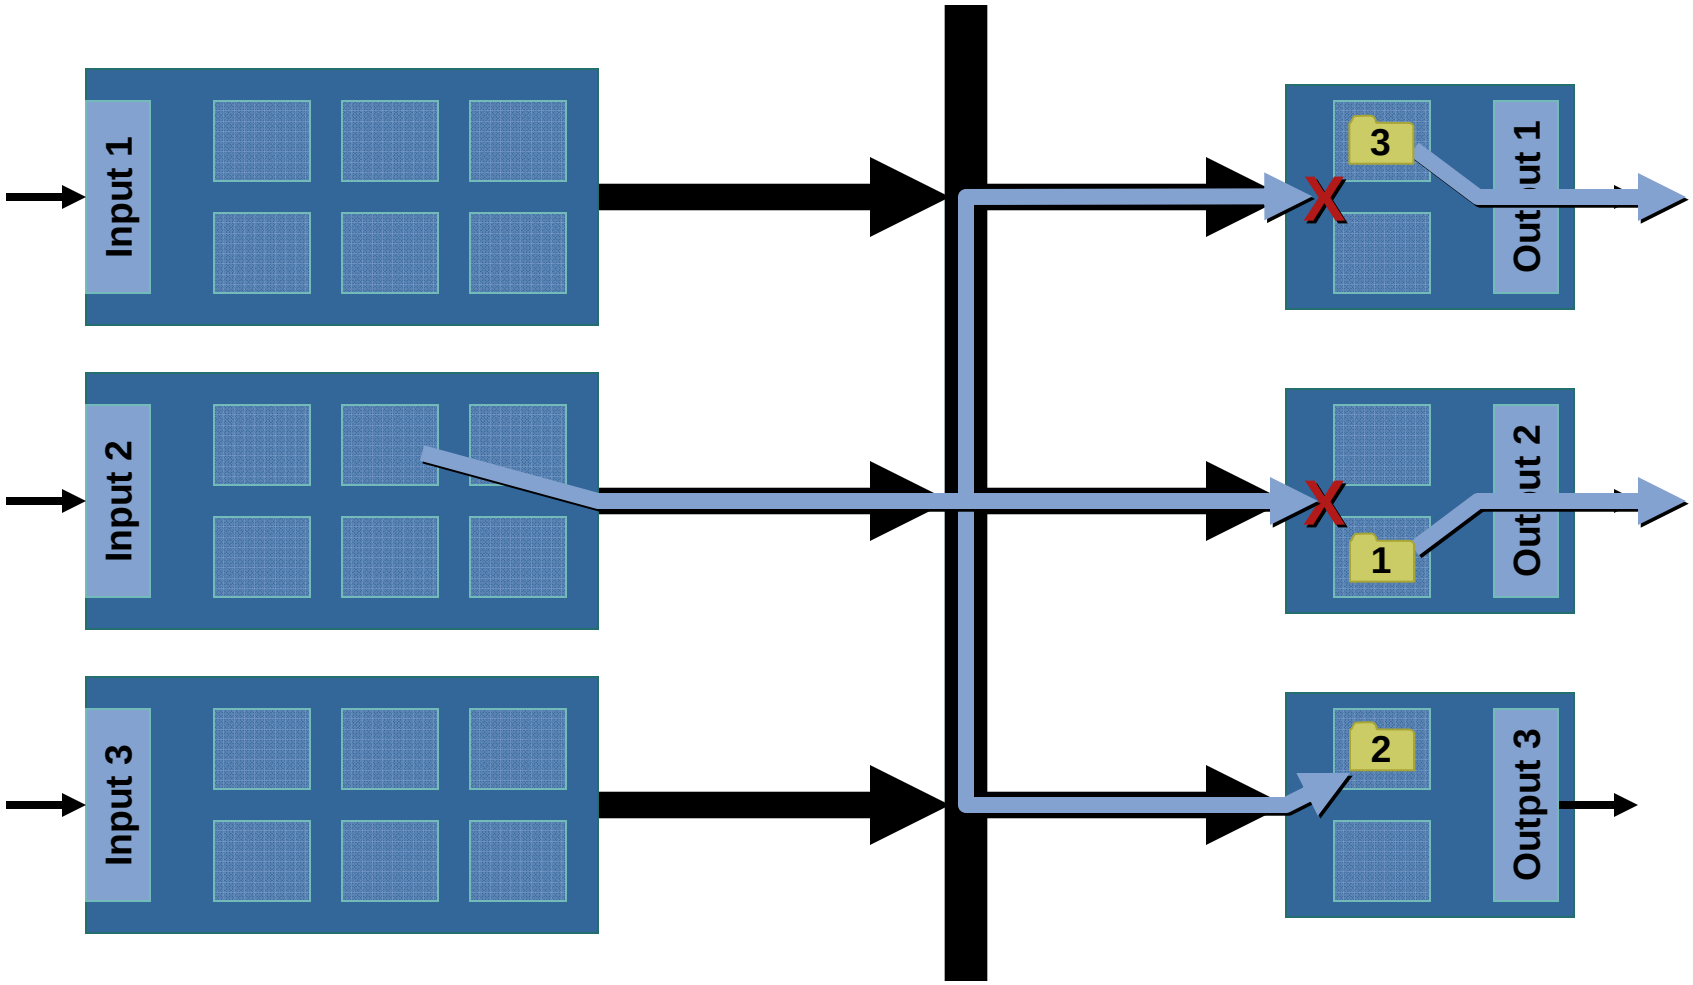

Bus

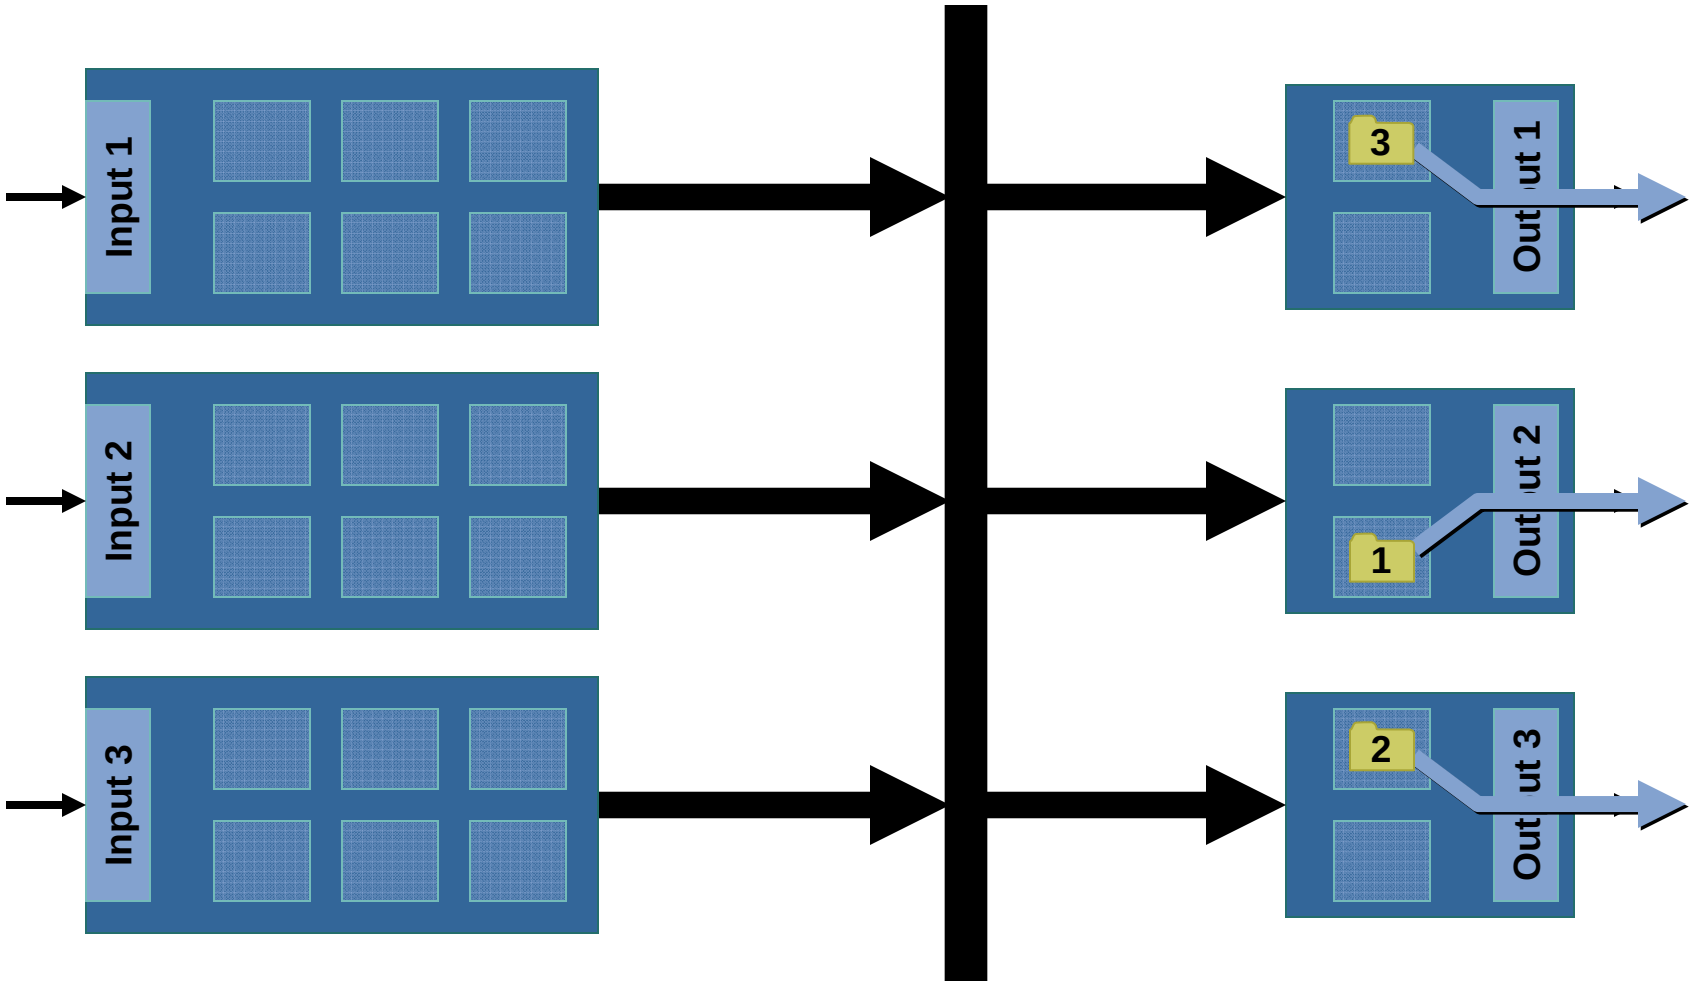

Bus

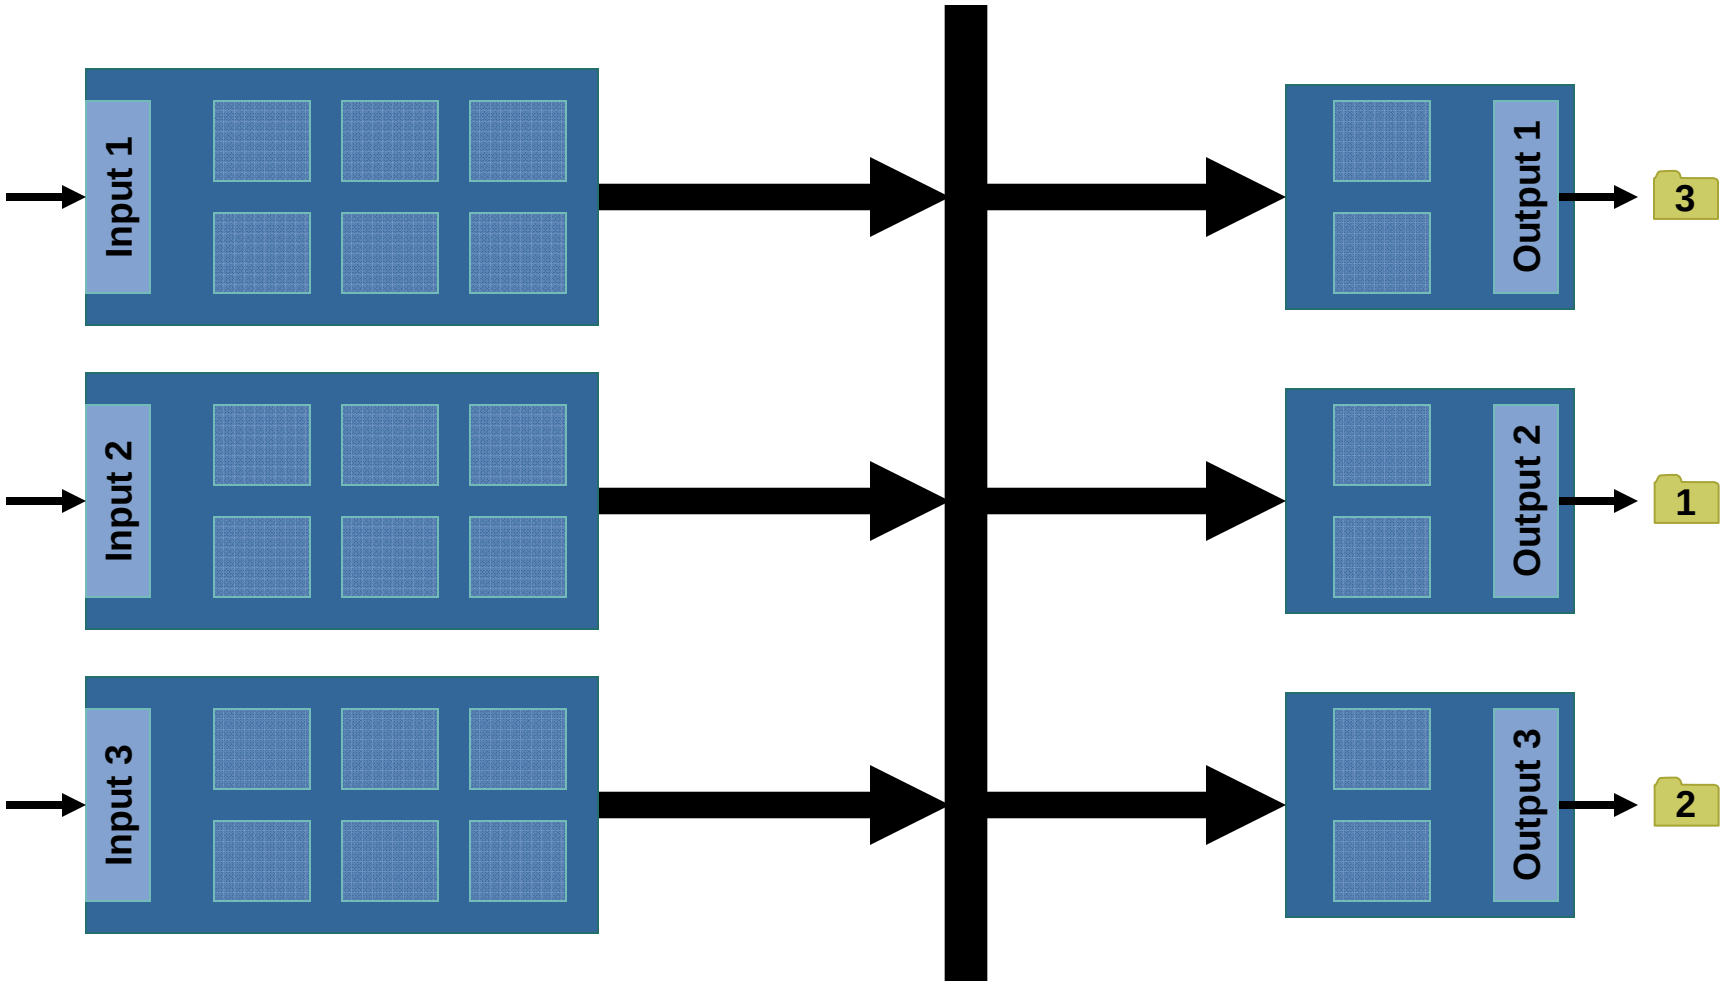

Bus

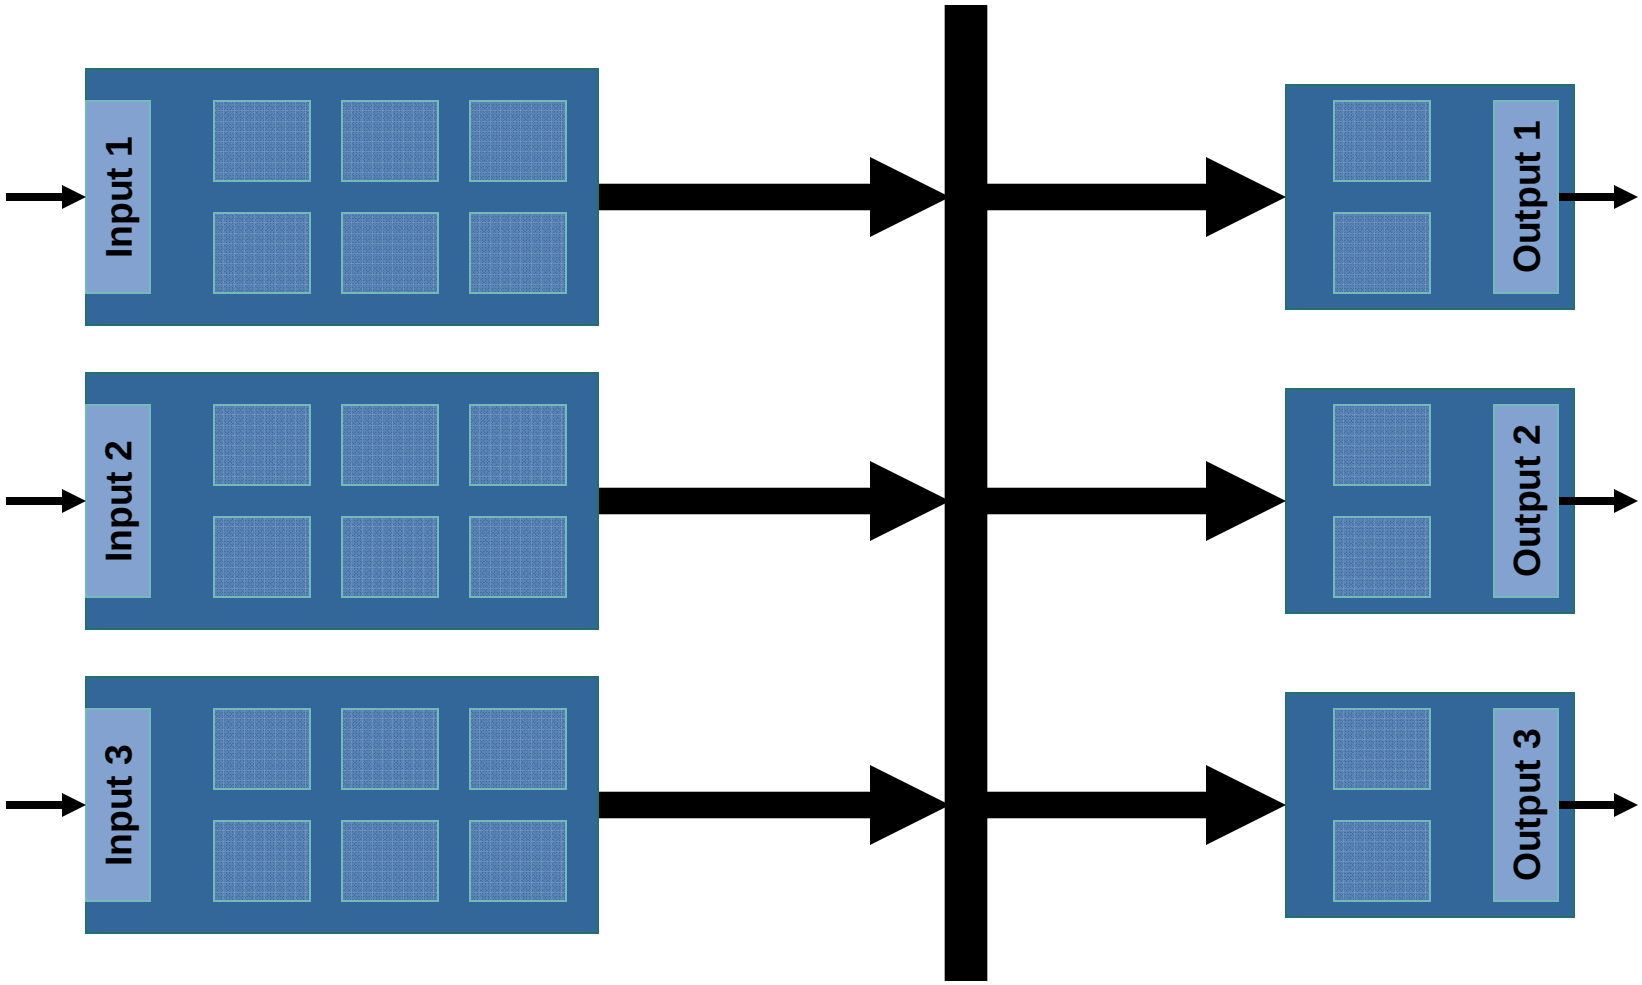

Bus with centralized memory

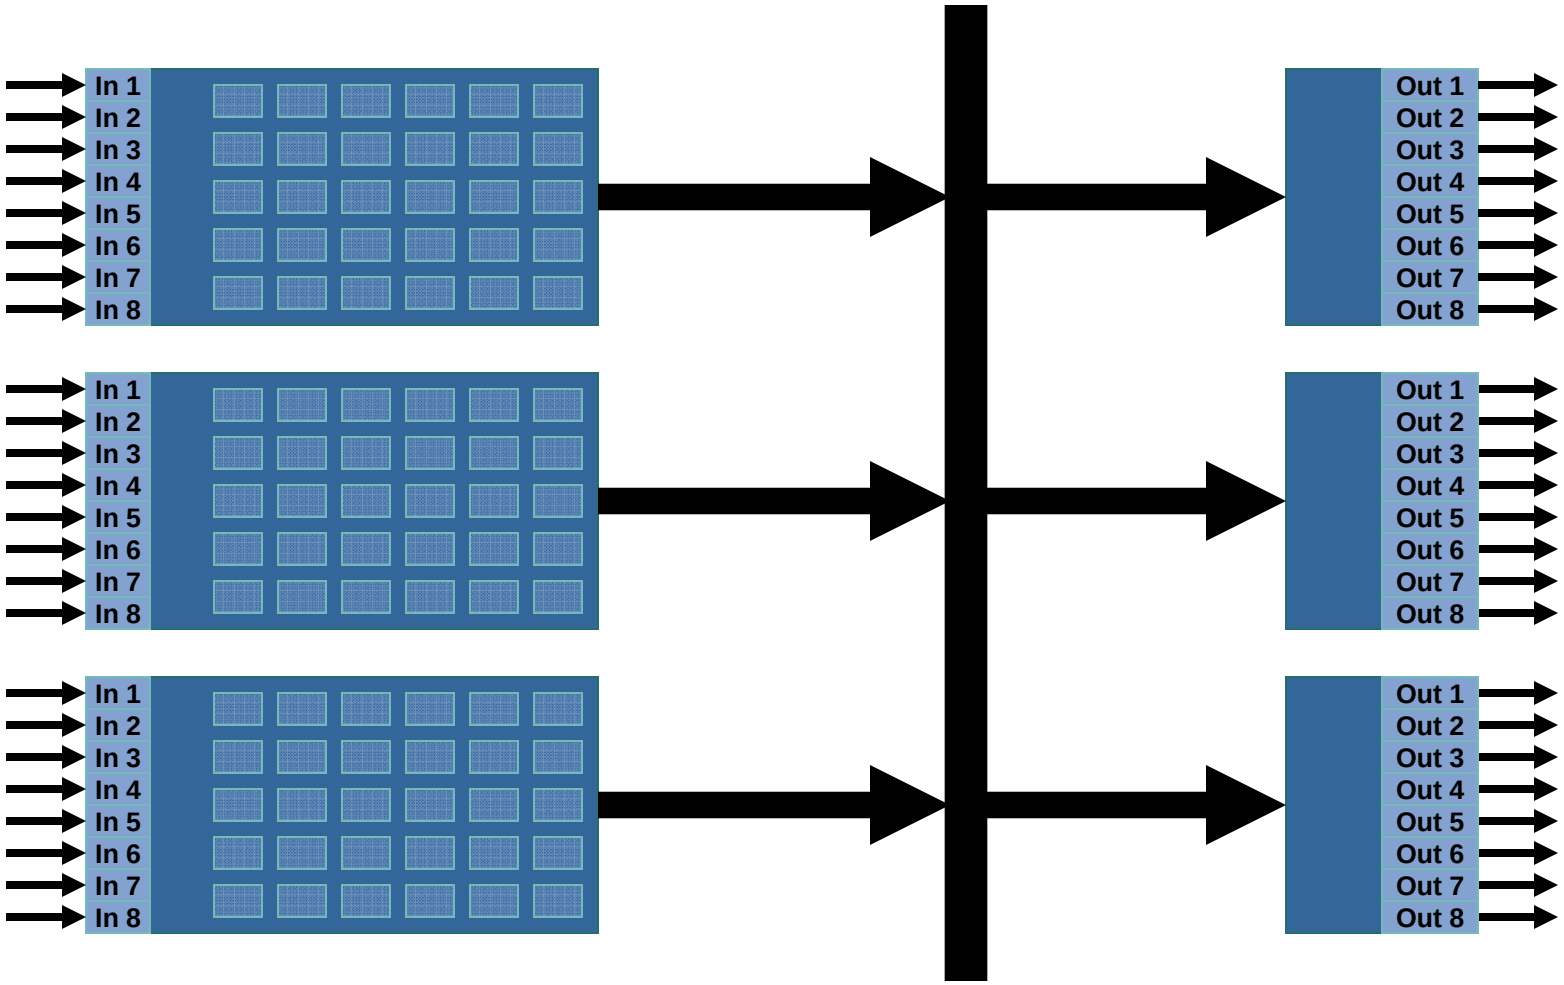

Part II—Architectures

- Centralized Memory
  - Switch on Chip / “System on Chip” (SoC)
- Bus
- **Mesh**
  - Centralized memory in a mesh
- Two-Tier / Banyan / Clos
  - Centralized memory in two tiers
- Cross-Bar

Mesh

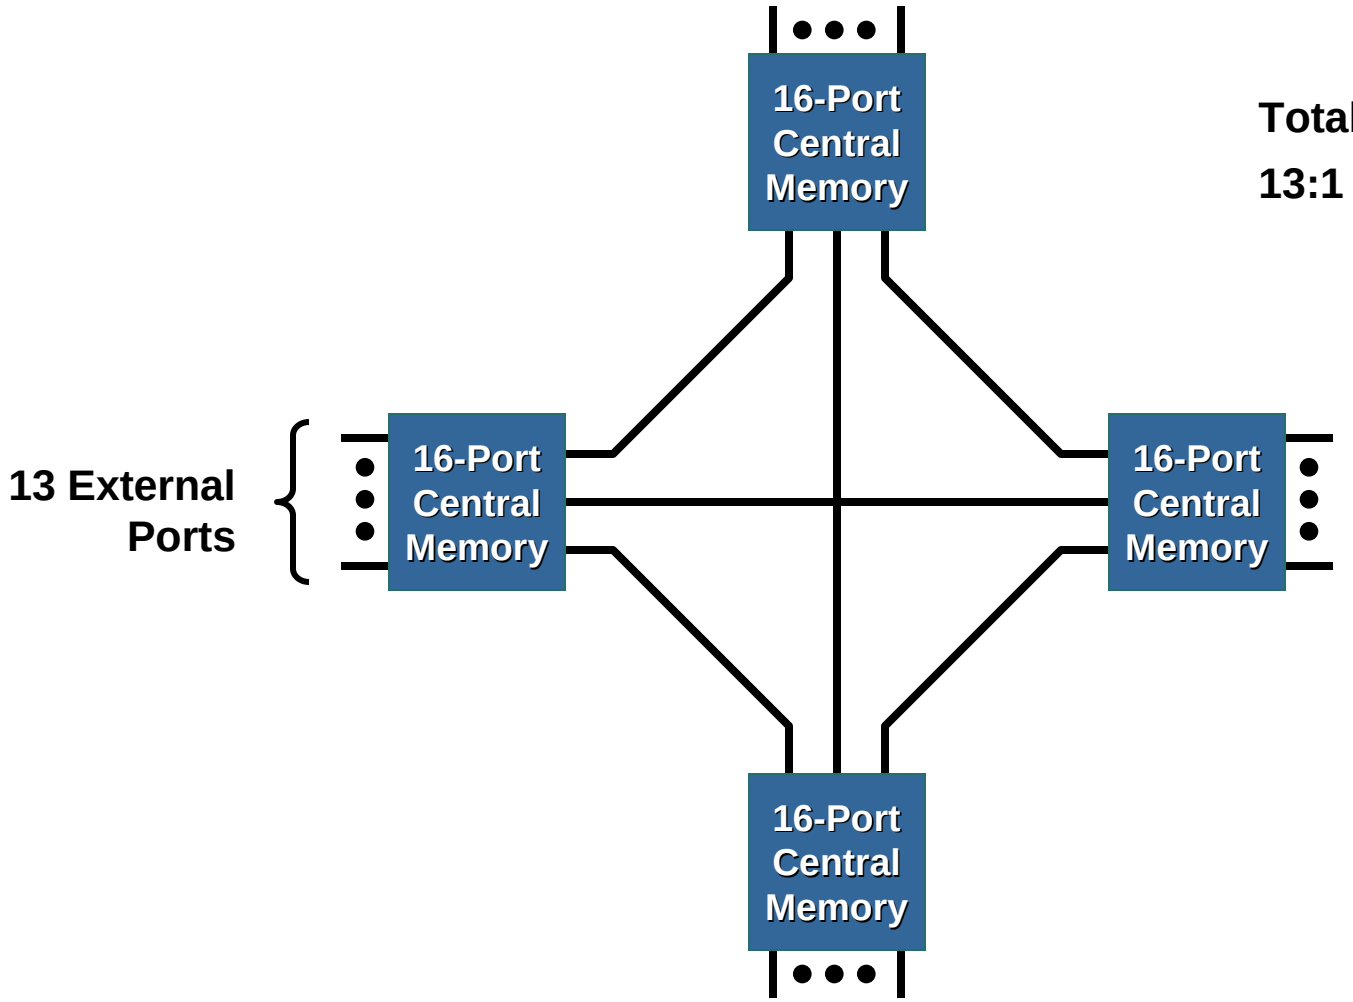

Total 52 External Ports

13:1 Blocking Factor

Mesh

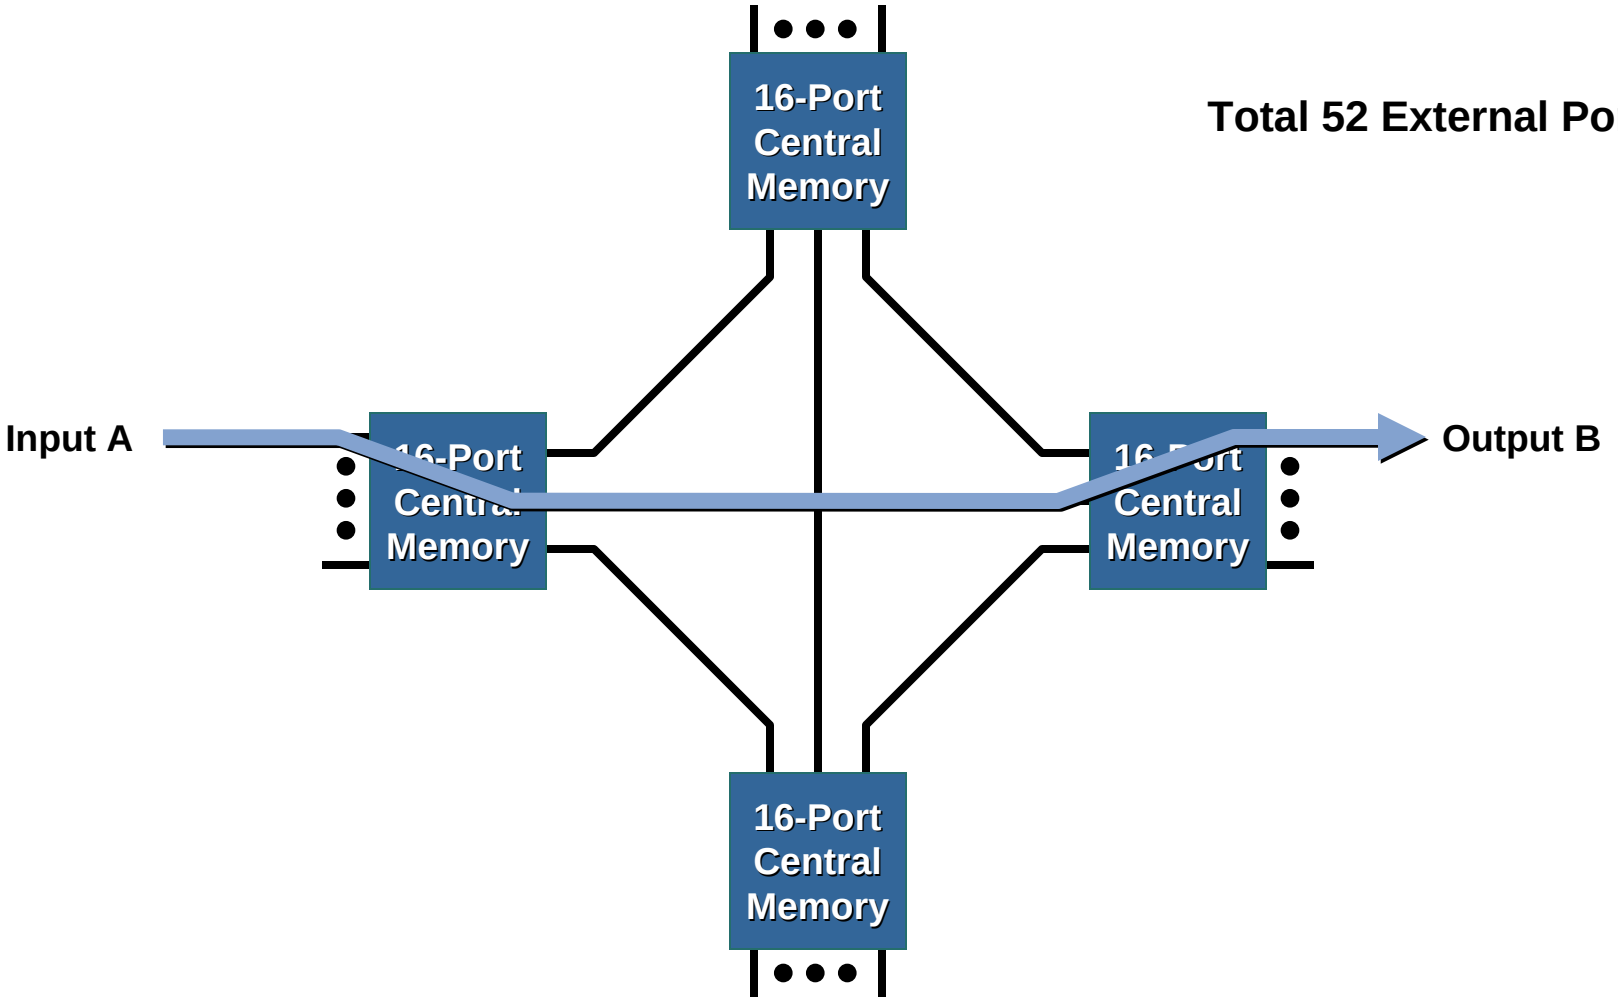

Mesh

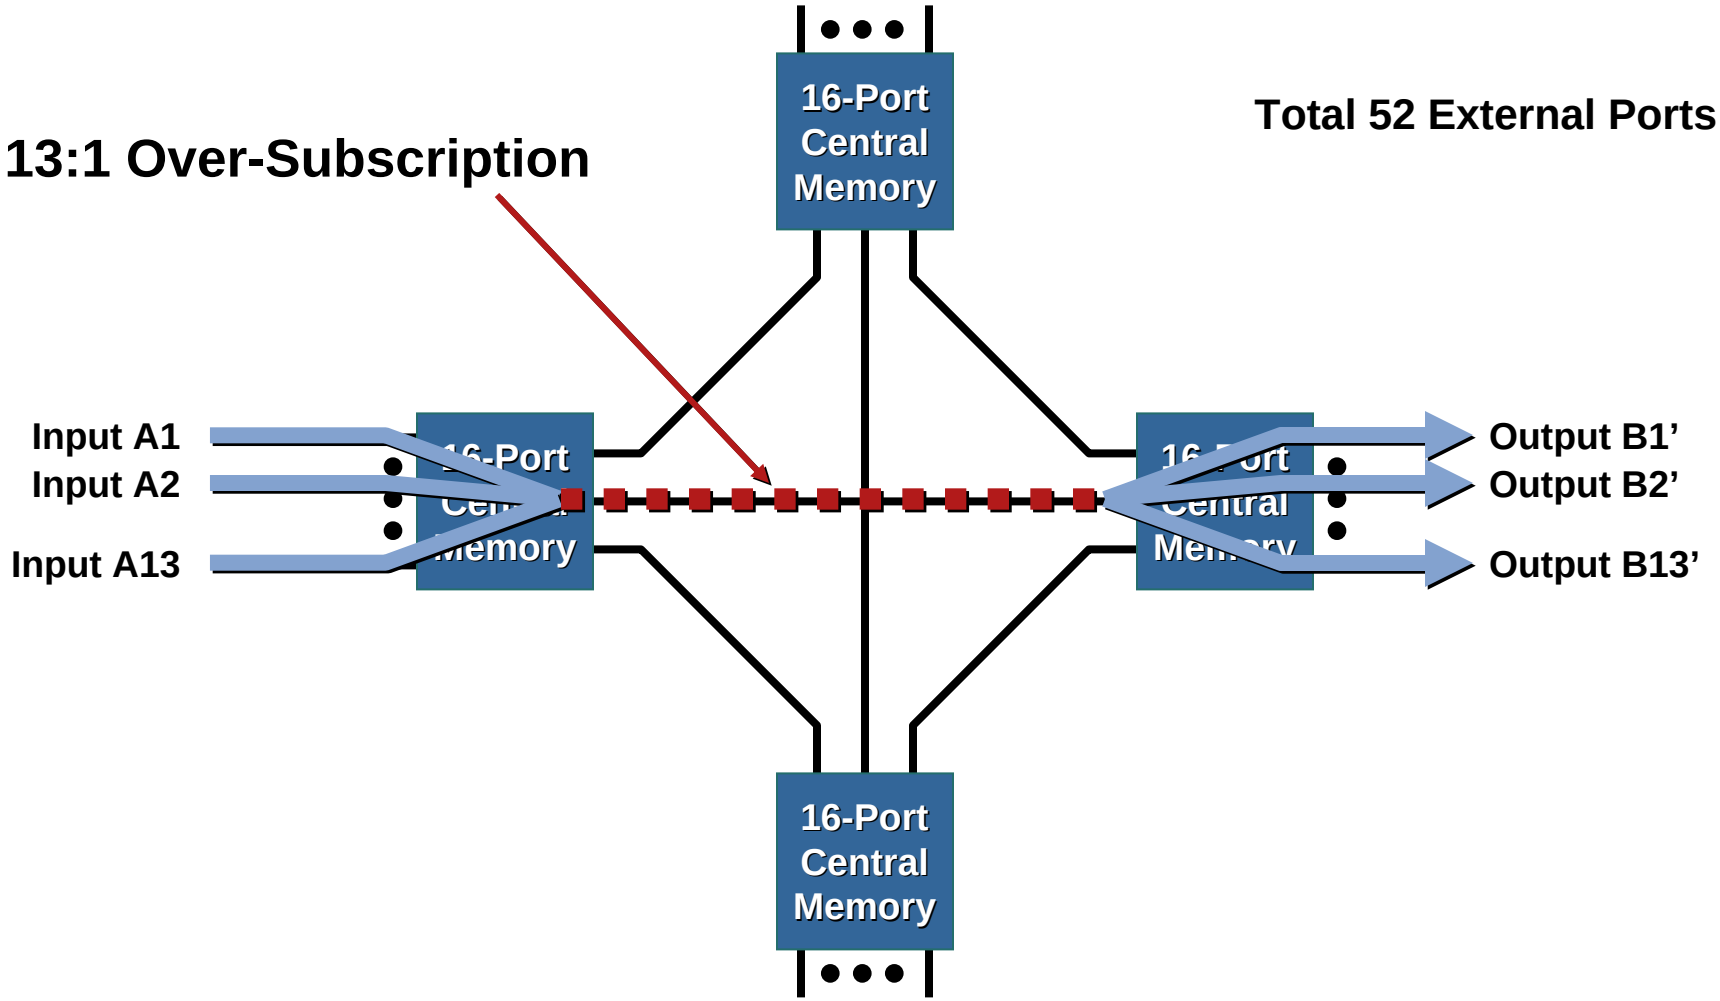

Mesh

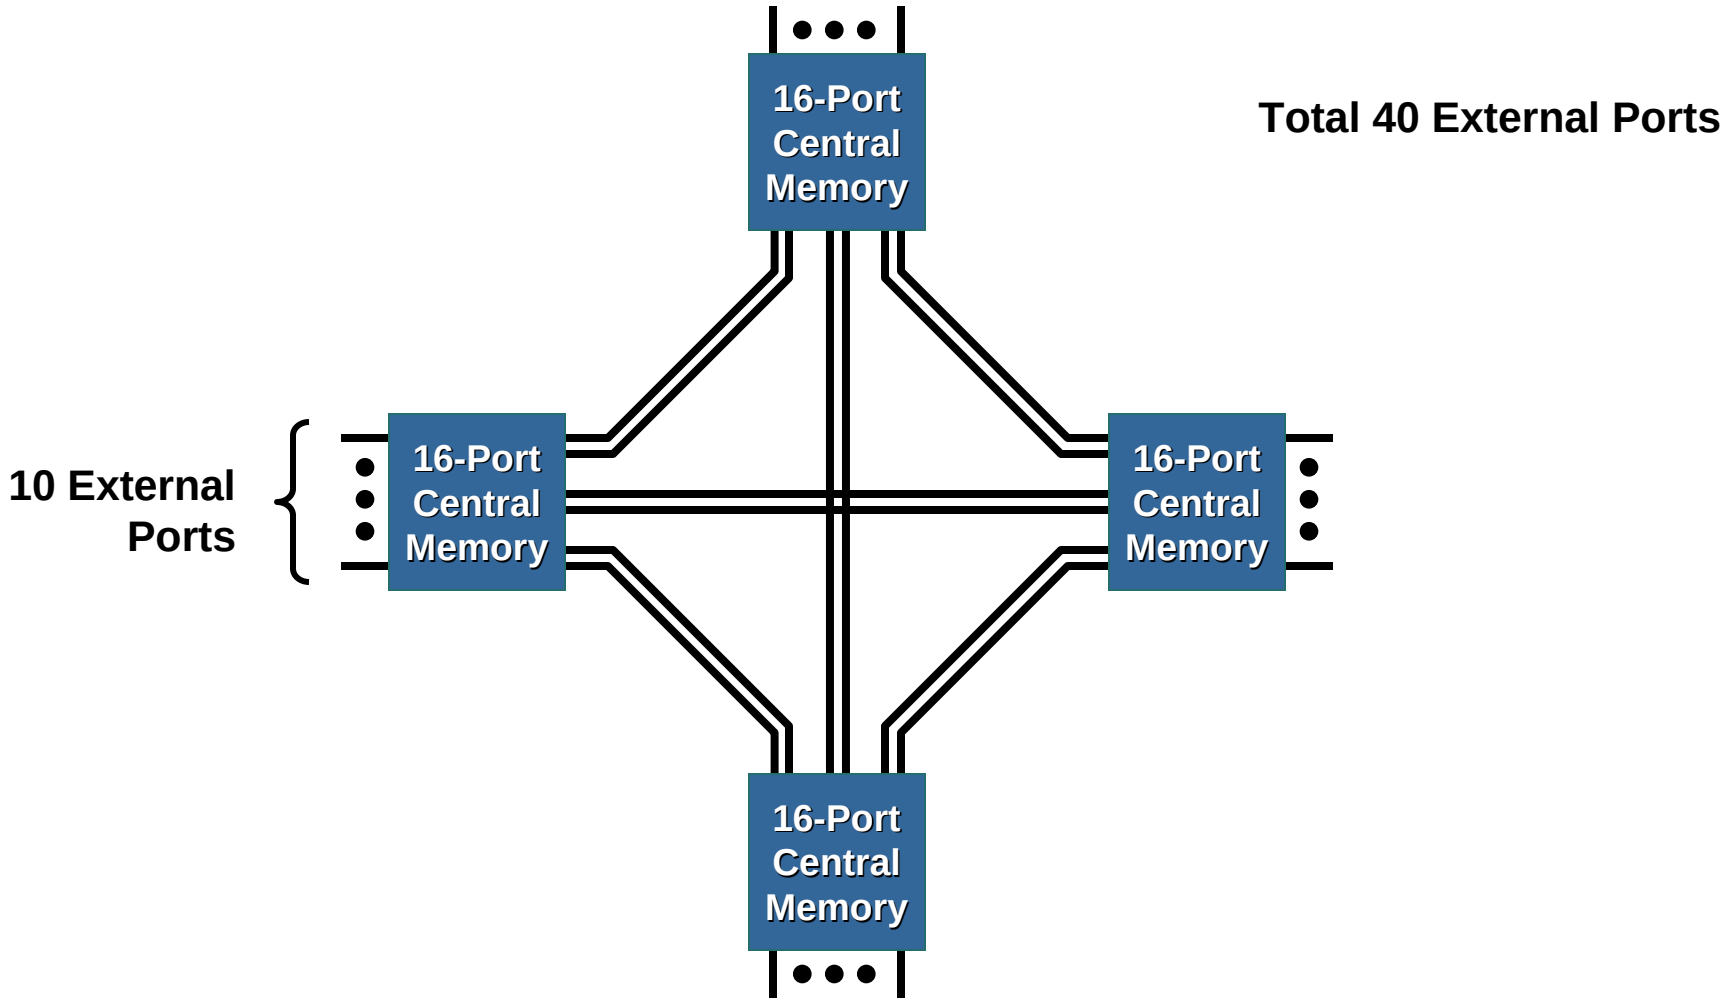

Mesh

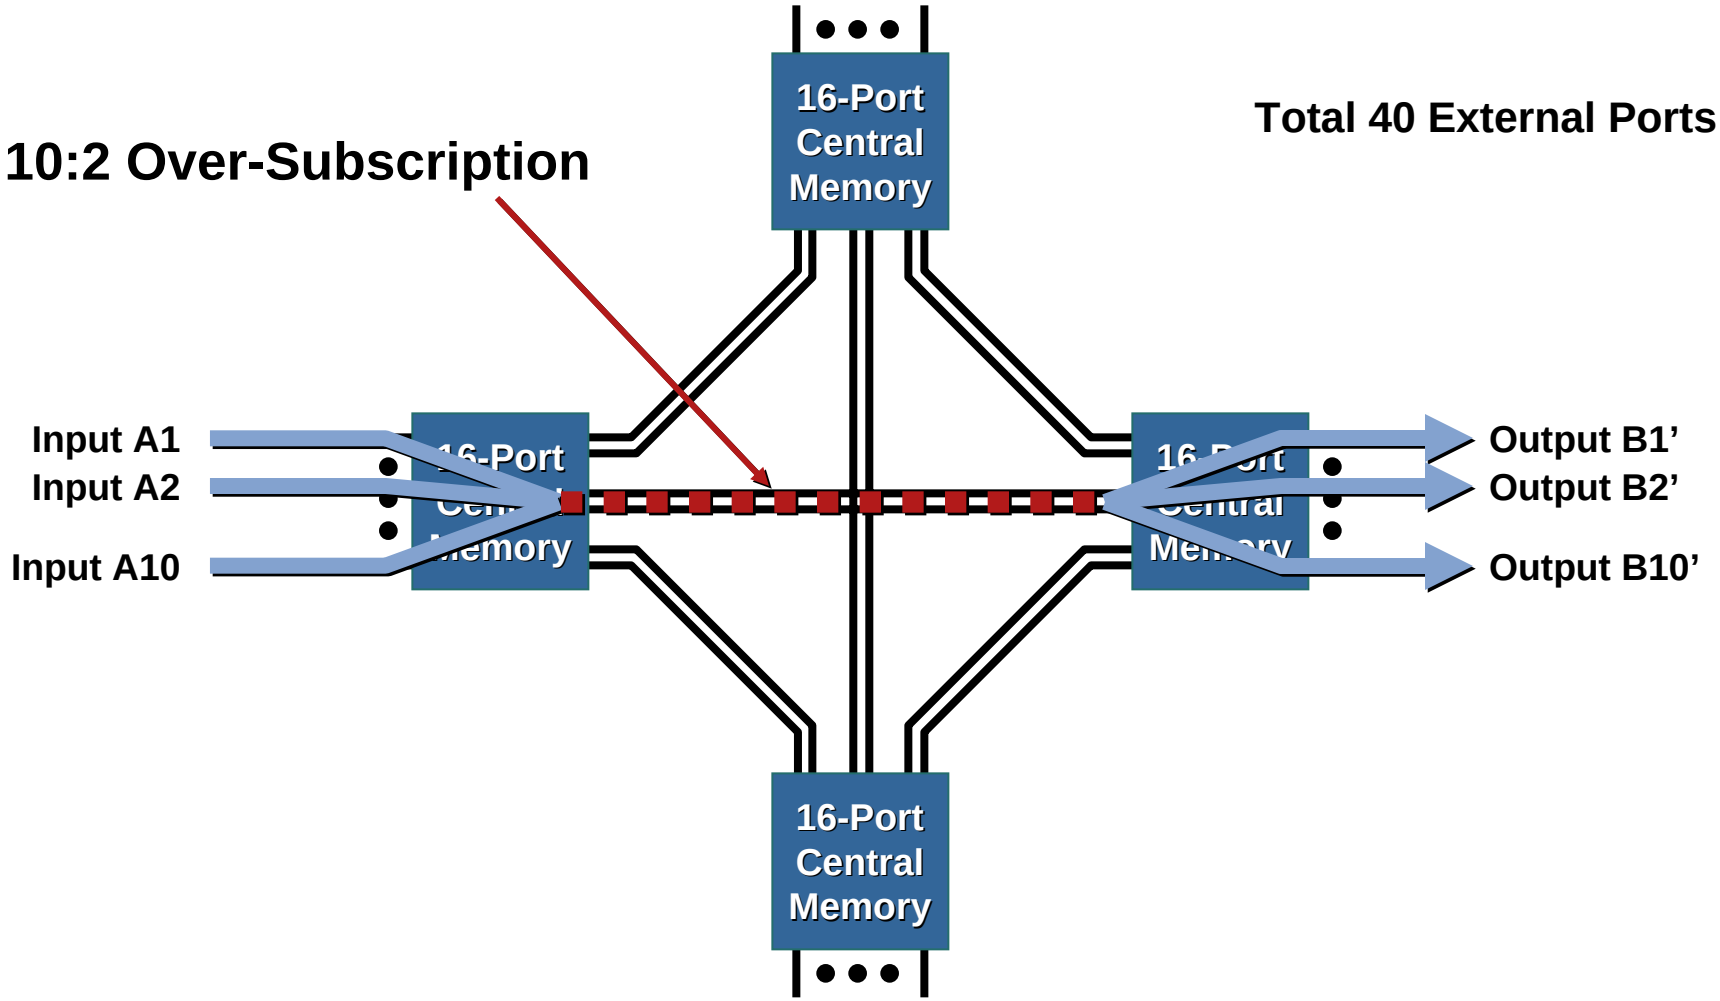

Mesh

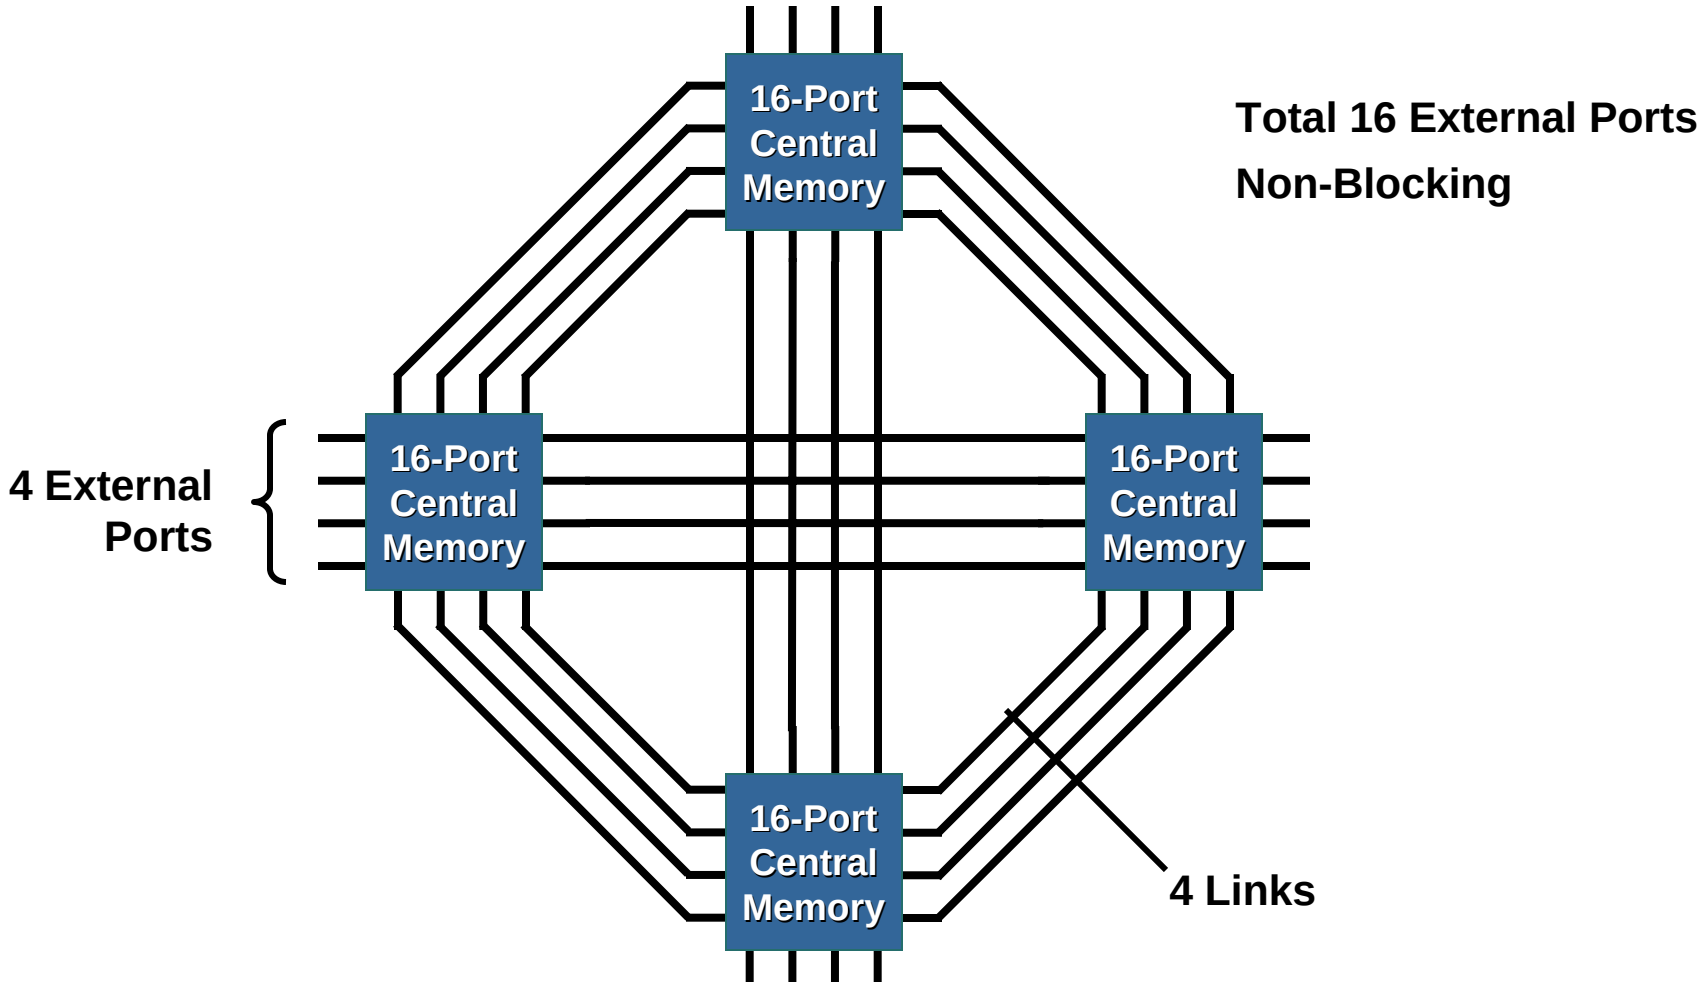

Mesh

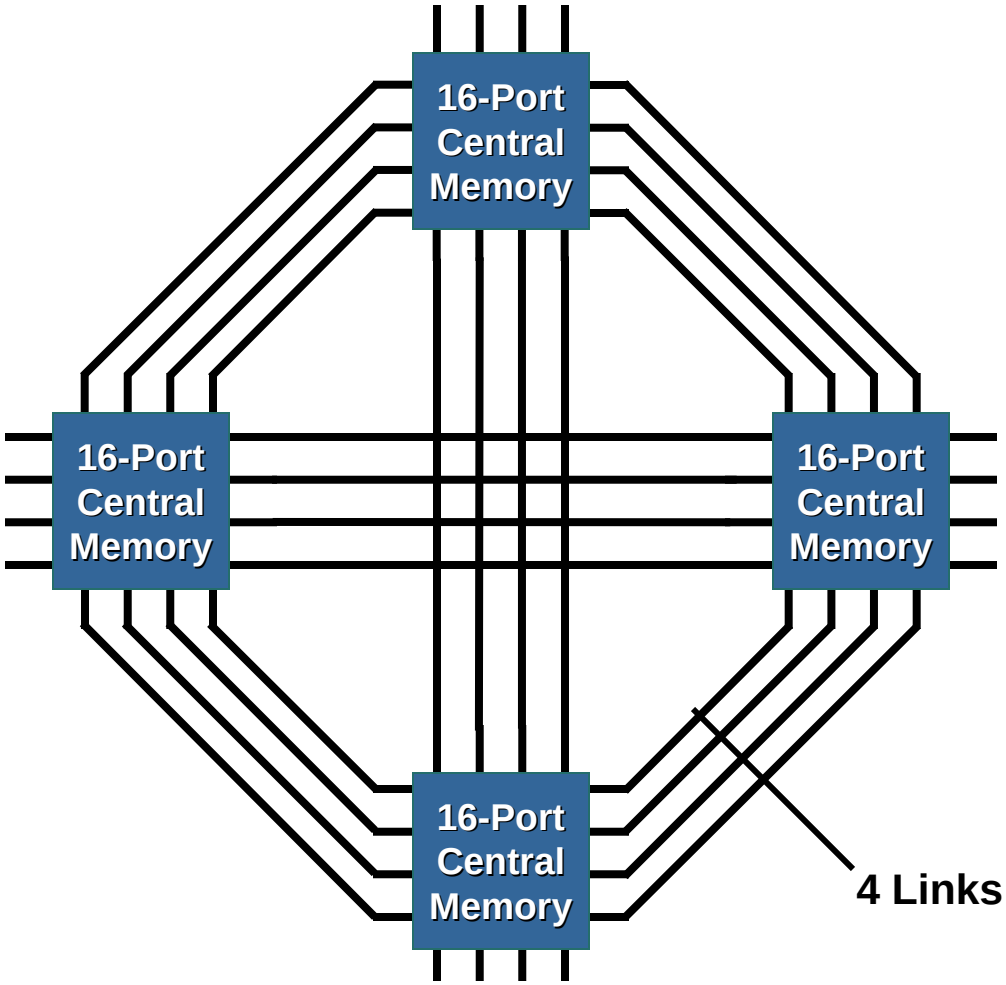

Total 16 External Ports

Non-Blocking

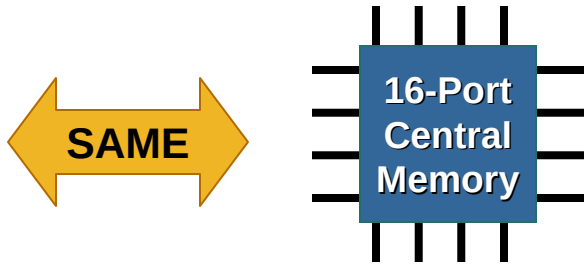

Mesh-Based Switch

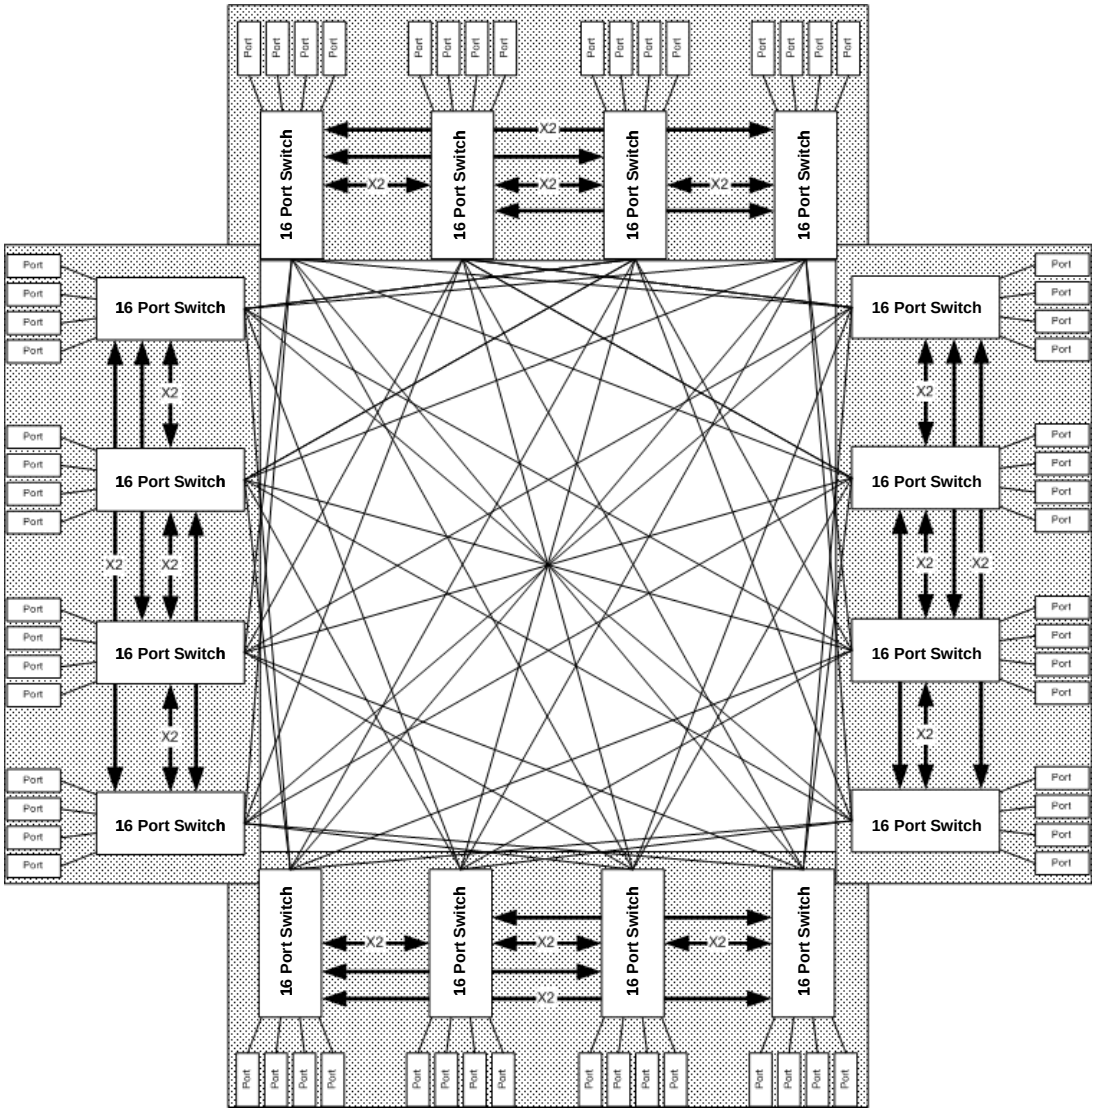

Mesh-Based Switch

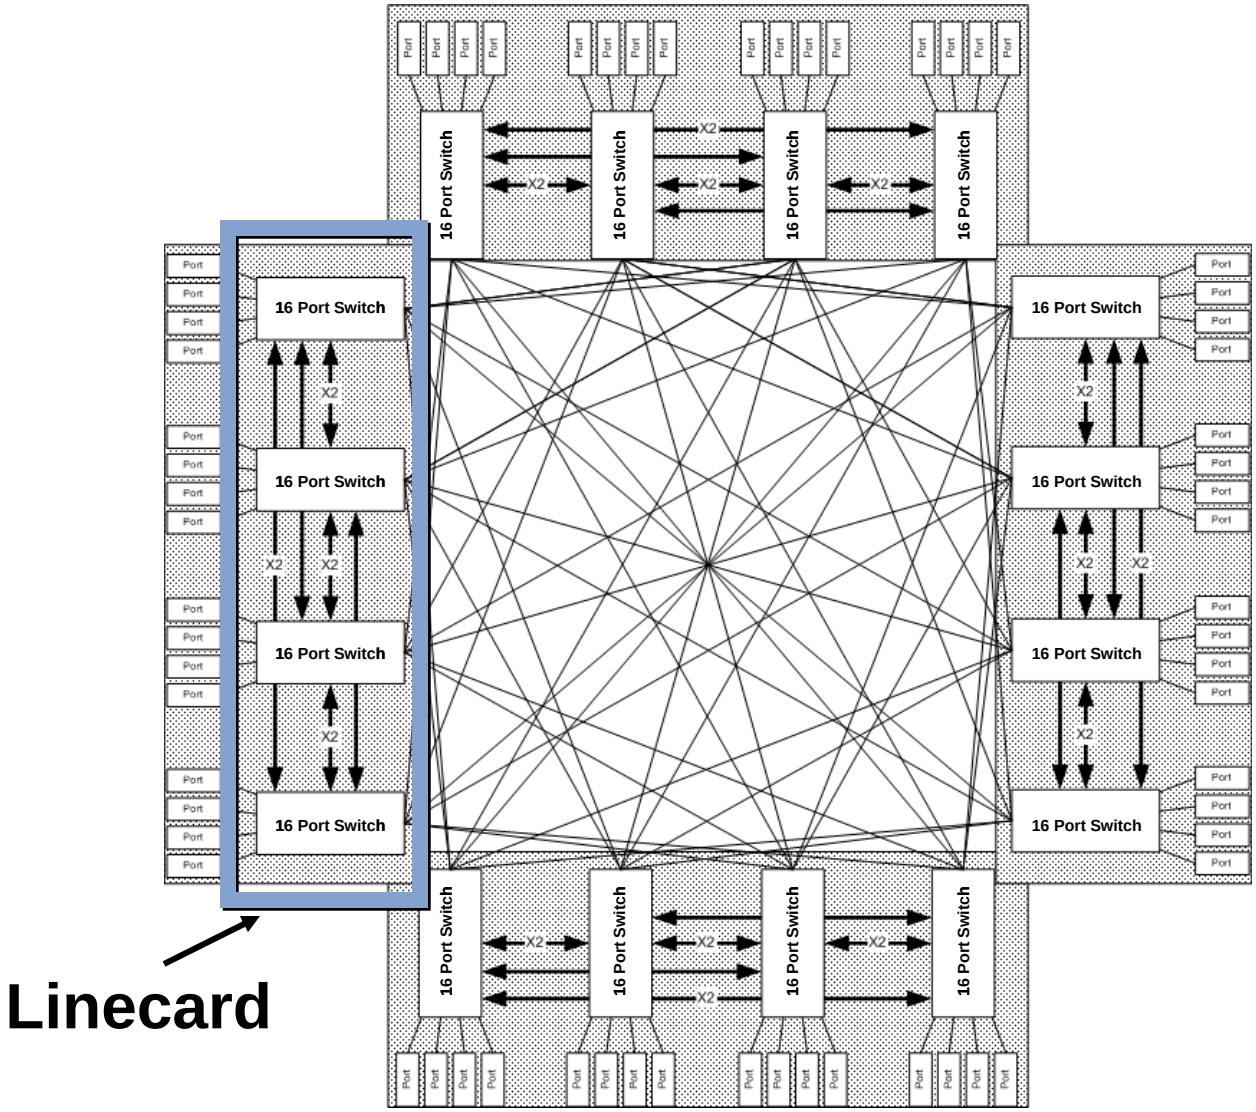

Mesh-Based Switch

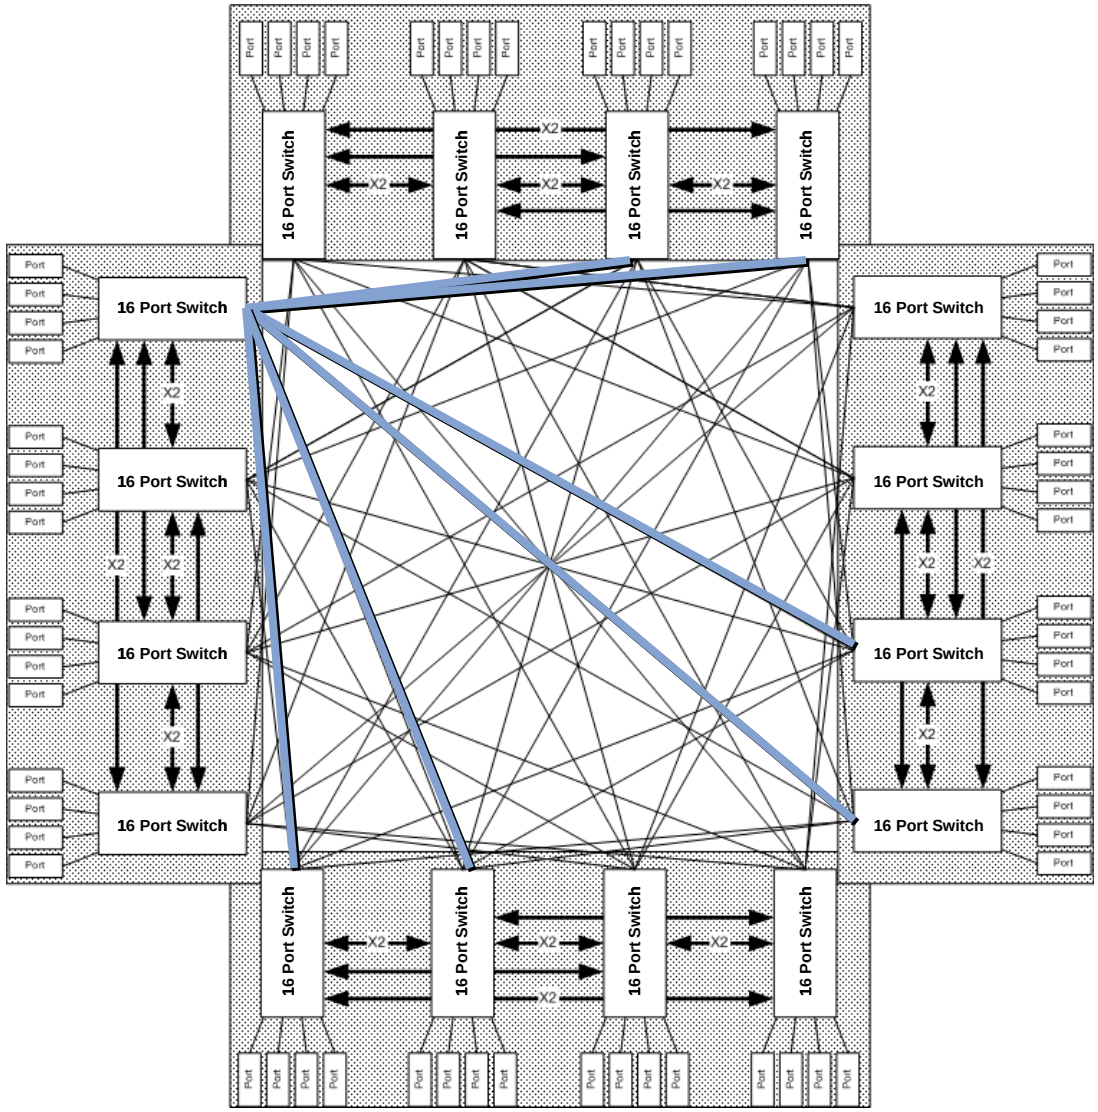

Mesh-Based Switch

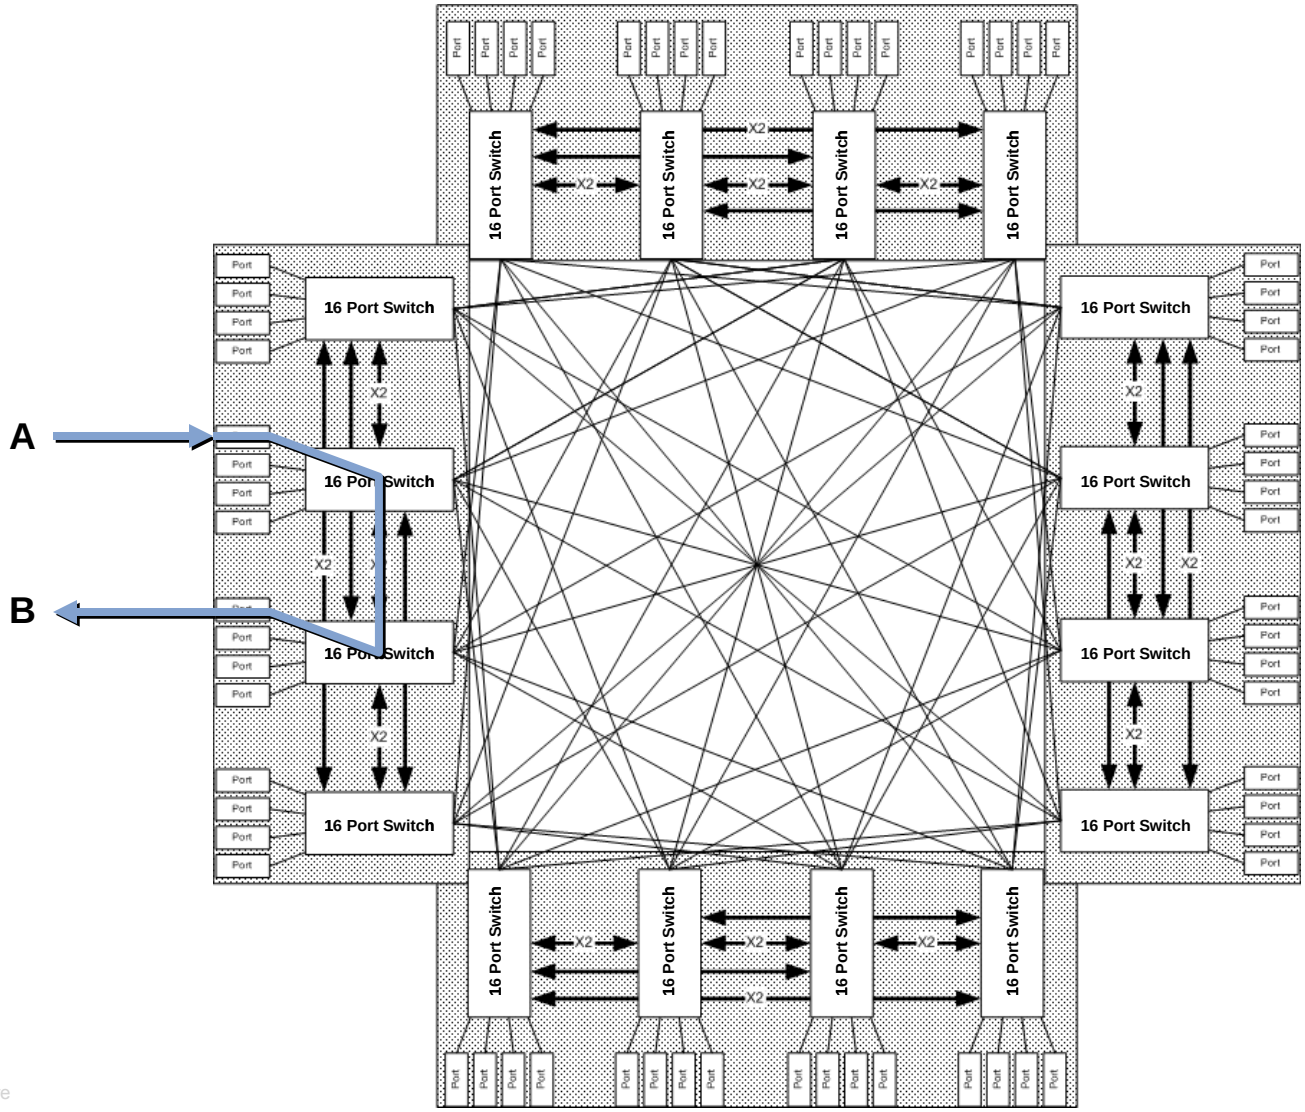

Mesh-Based Switch

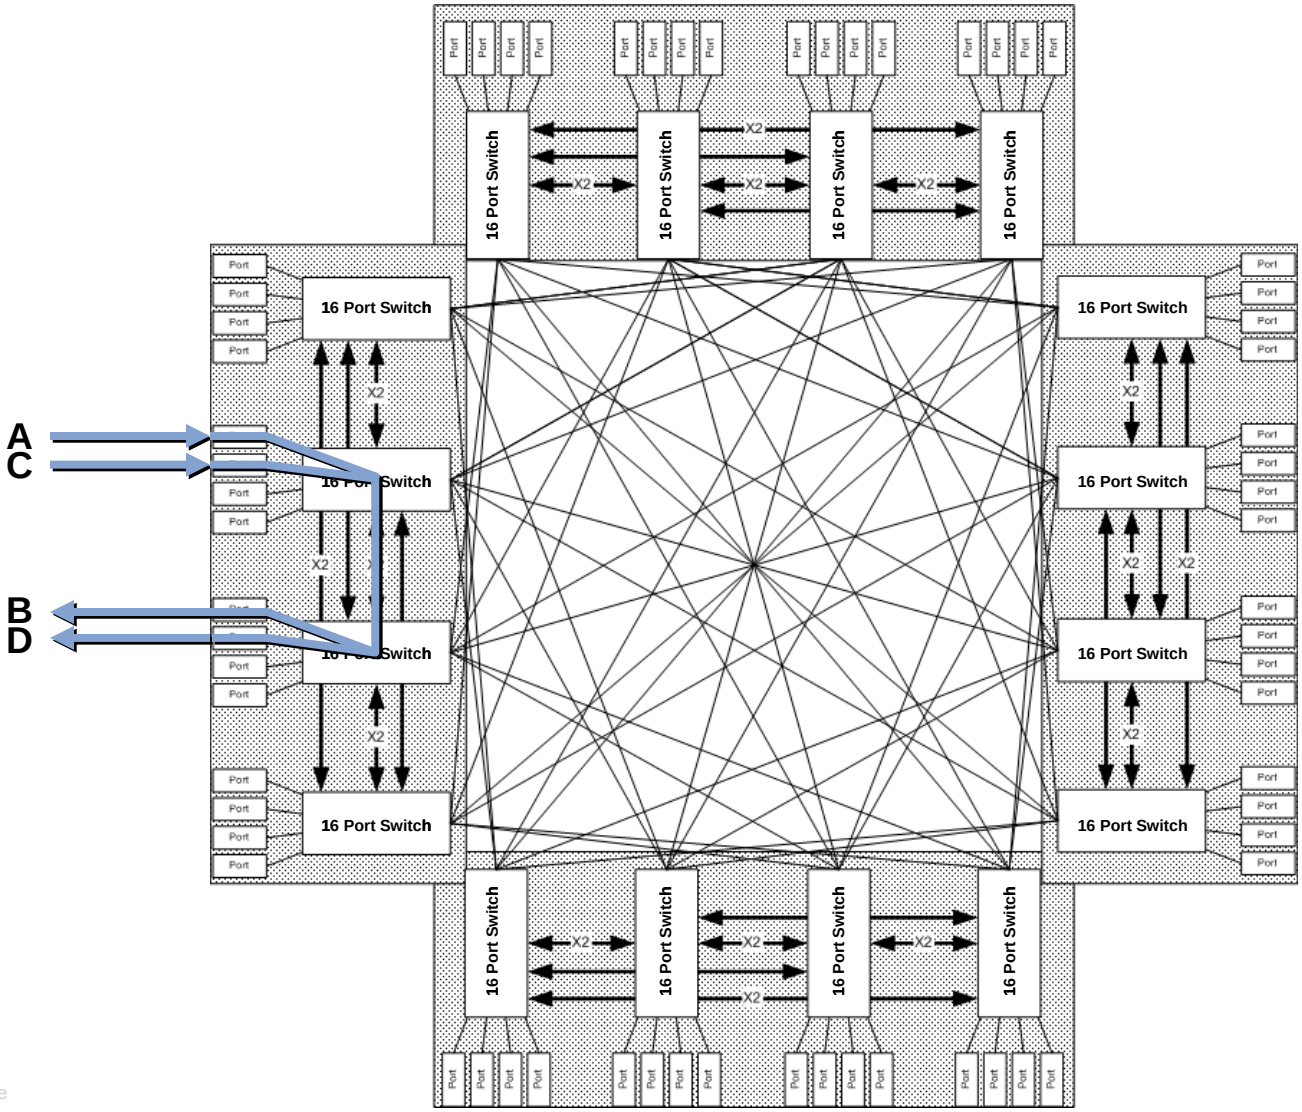

Mesh Based Switch

2:1 Over-Subscription

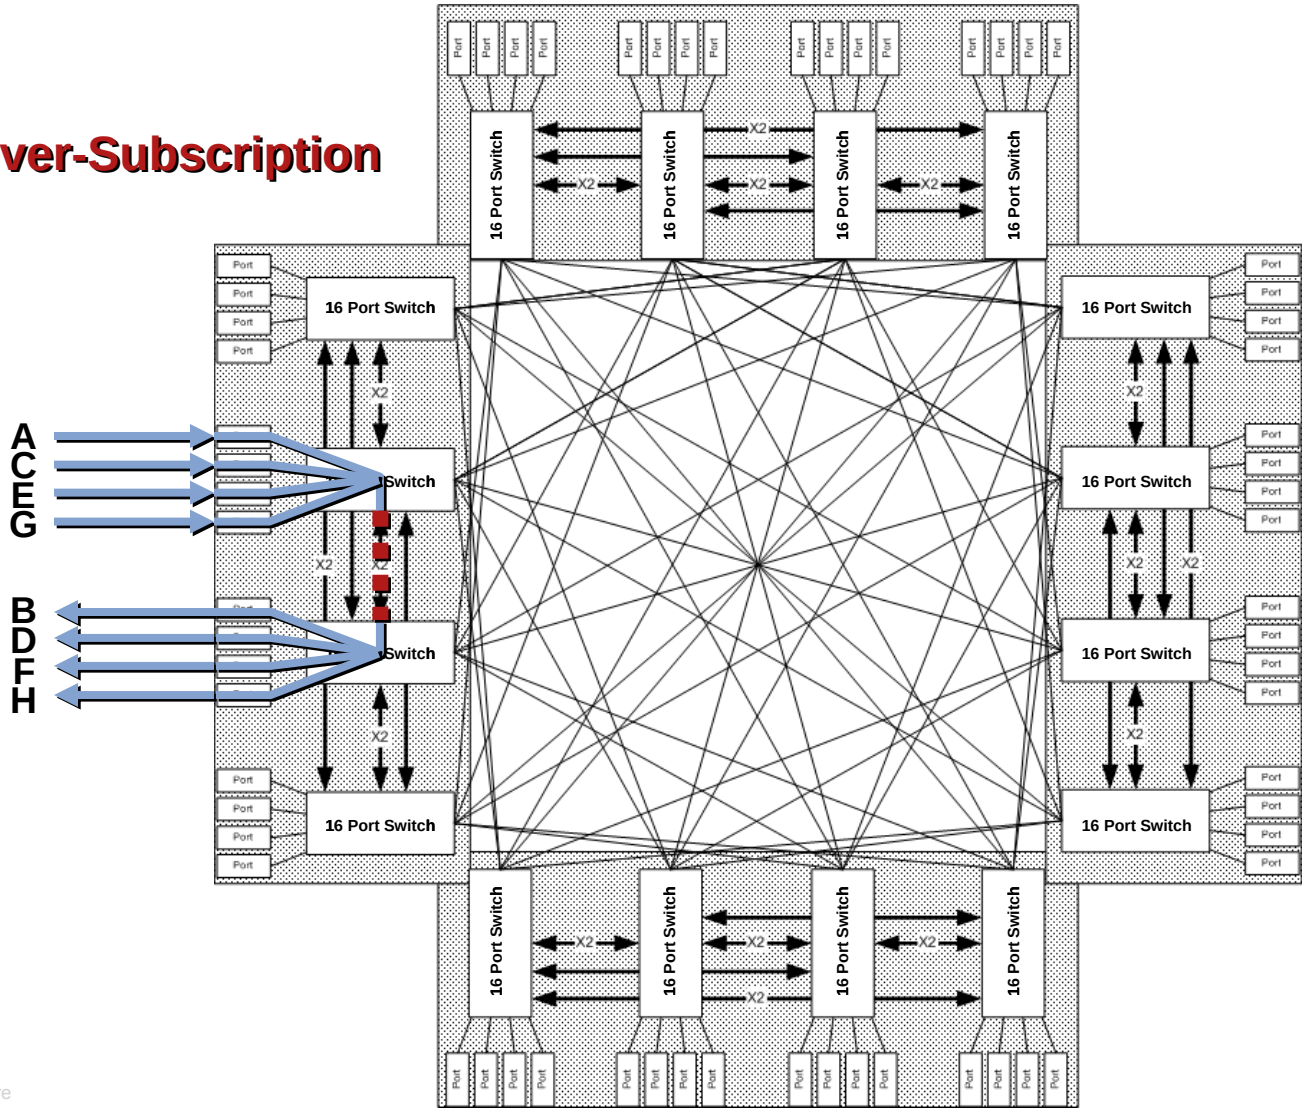

Mesh-Based Switch

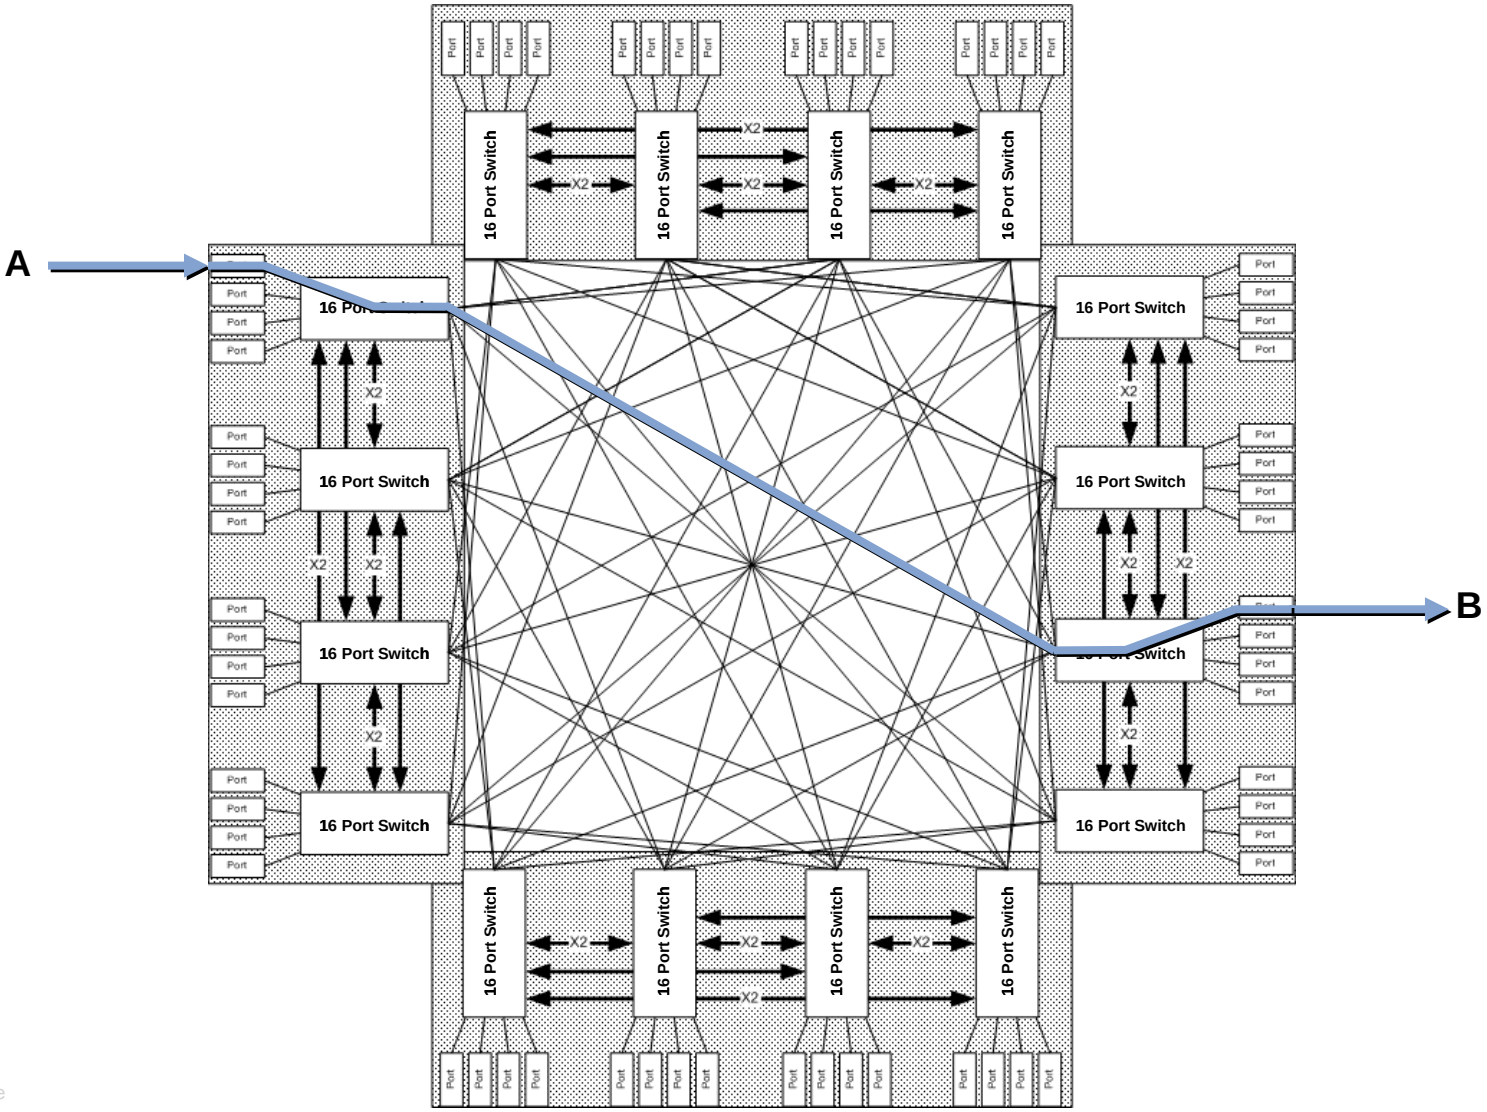

Mesh-Based Switch

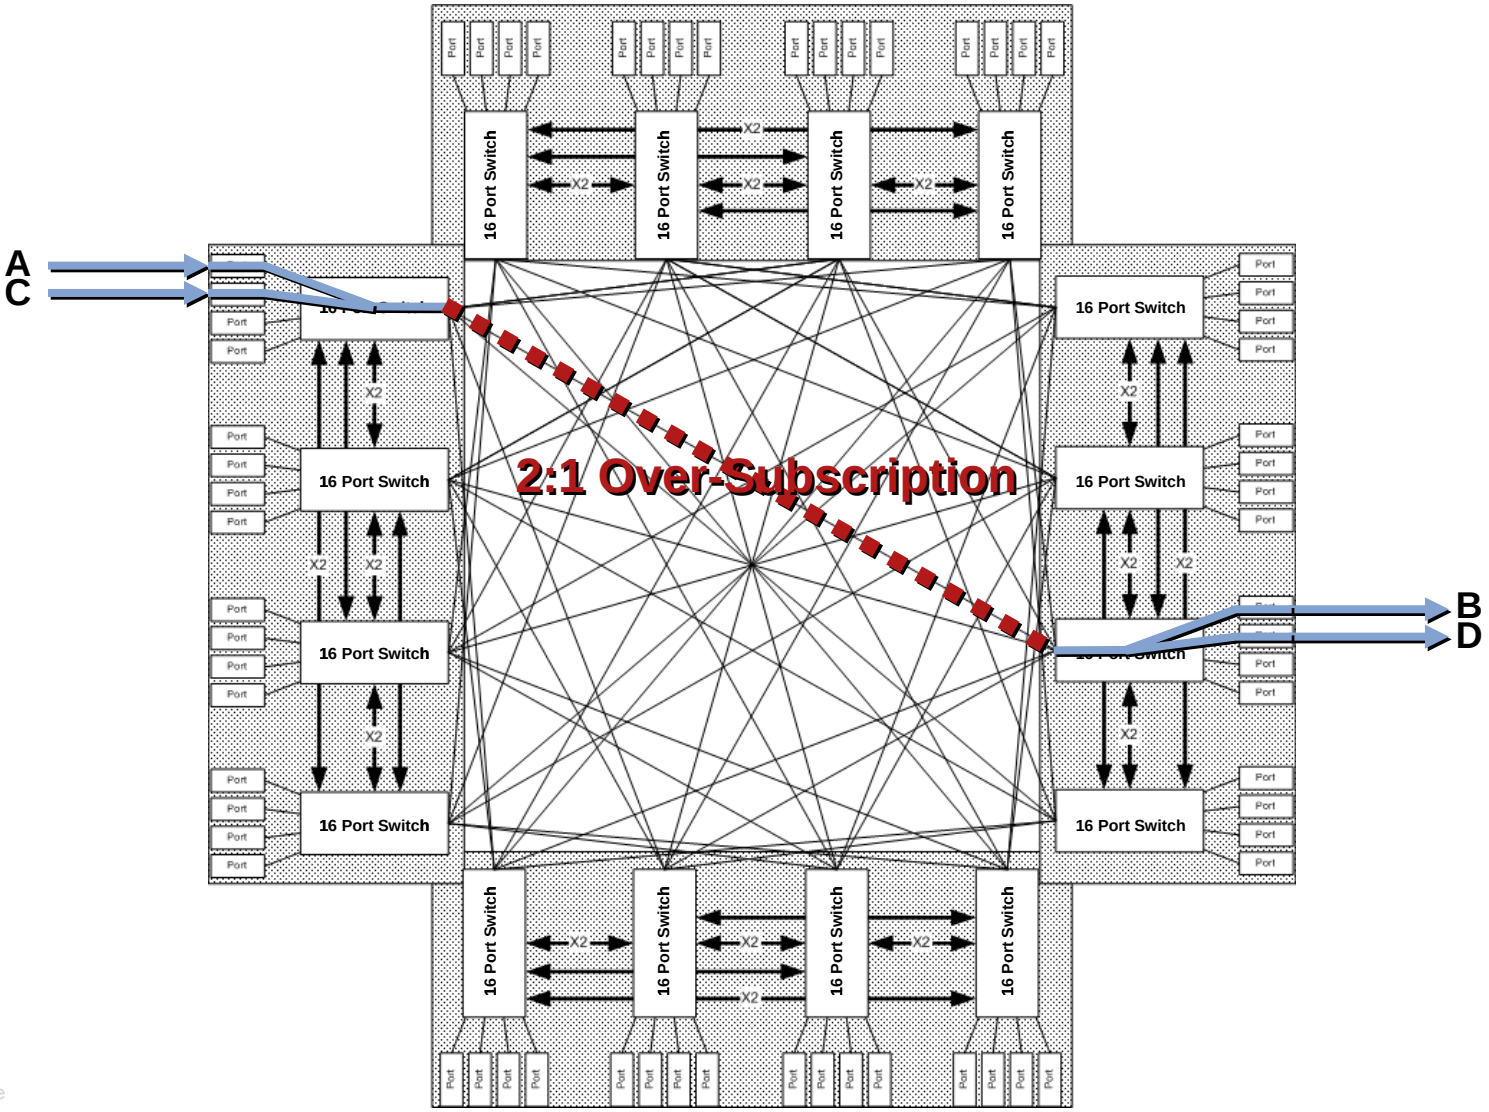

Mesh-Based Switch

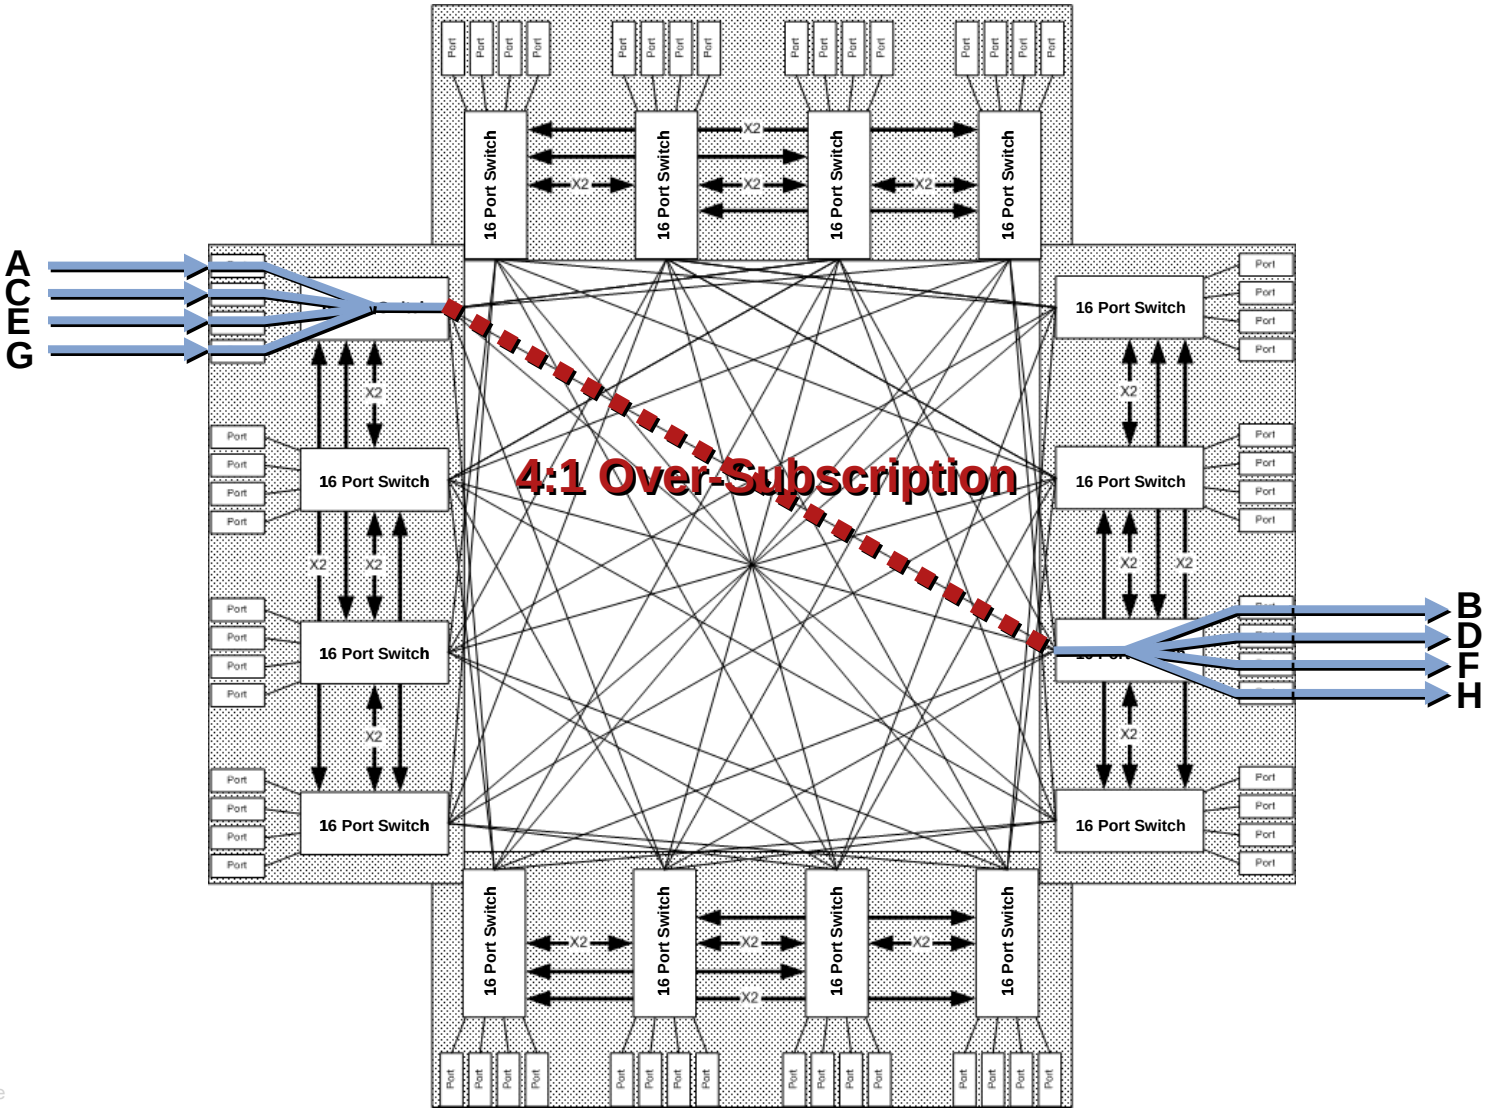

Mesh-Based Switch

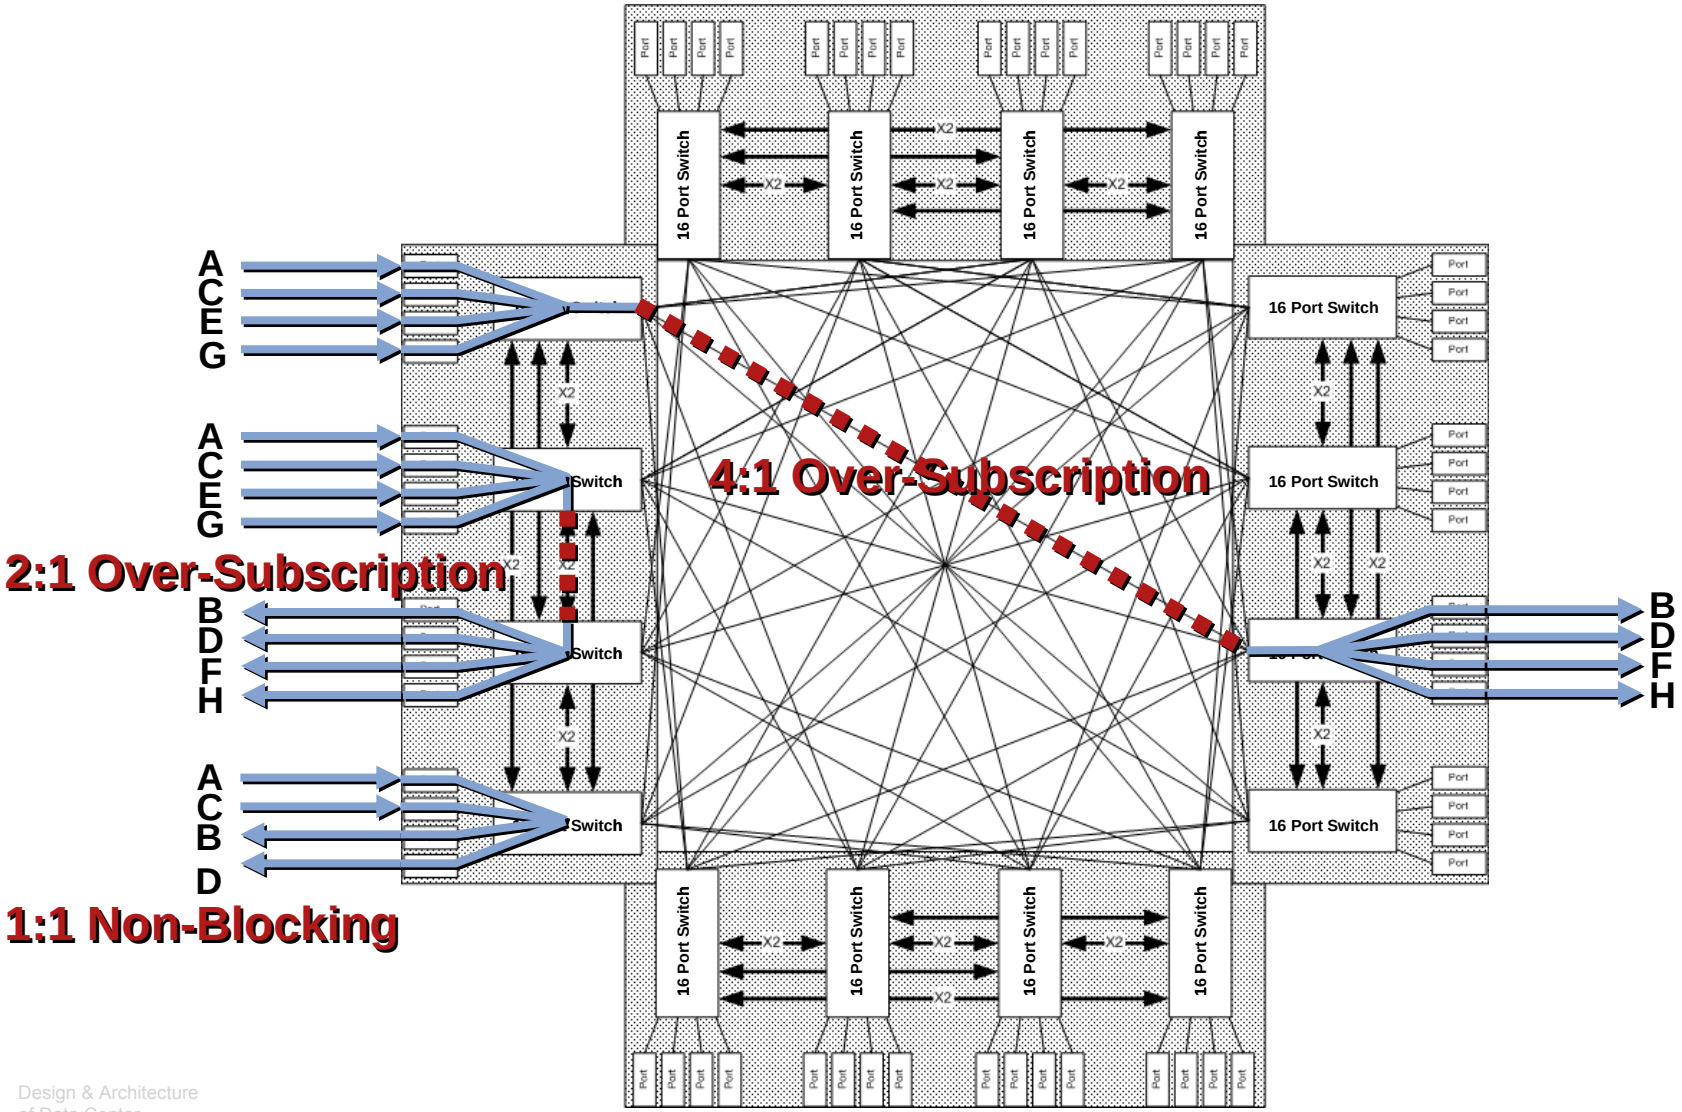

Part II—Architectures

- Centralized Memory
  - Switch on Chip / “System on Chip” (SoC)
- Bus
- Mesh
  - Centralized memory in a mesh
- **Two-Tier / Banyan / Clos**
  - Centralized memory in two tiers
- Cross-Bar

Two-Tier a.k.a. Banyan a.k.a. Clos

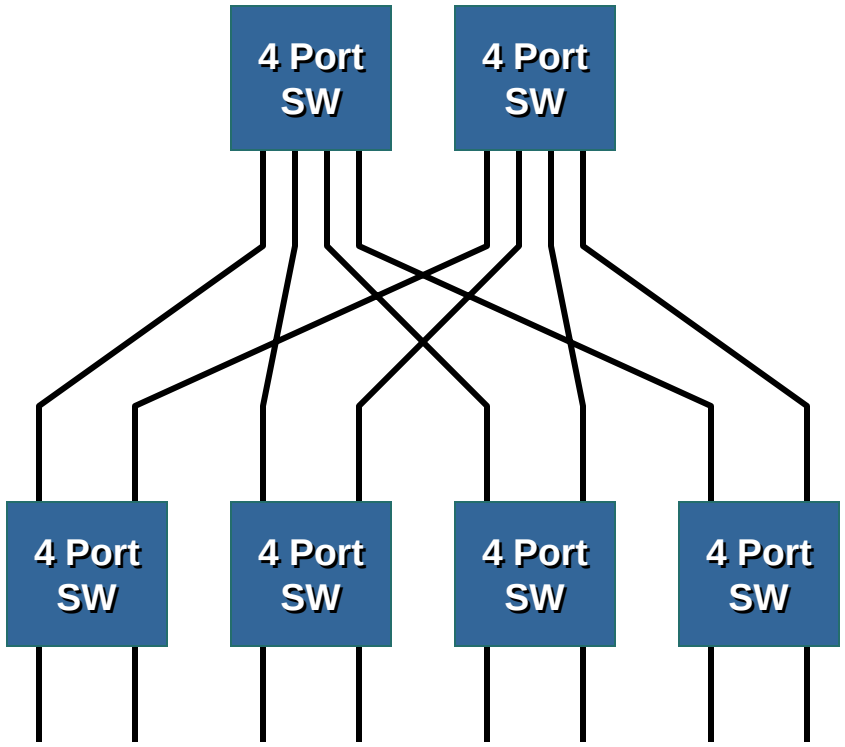

8 Port Switch

Two-Tier

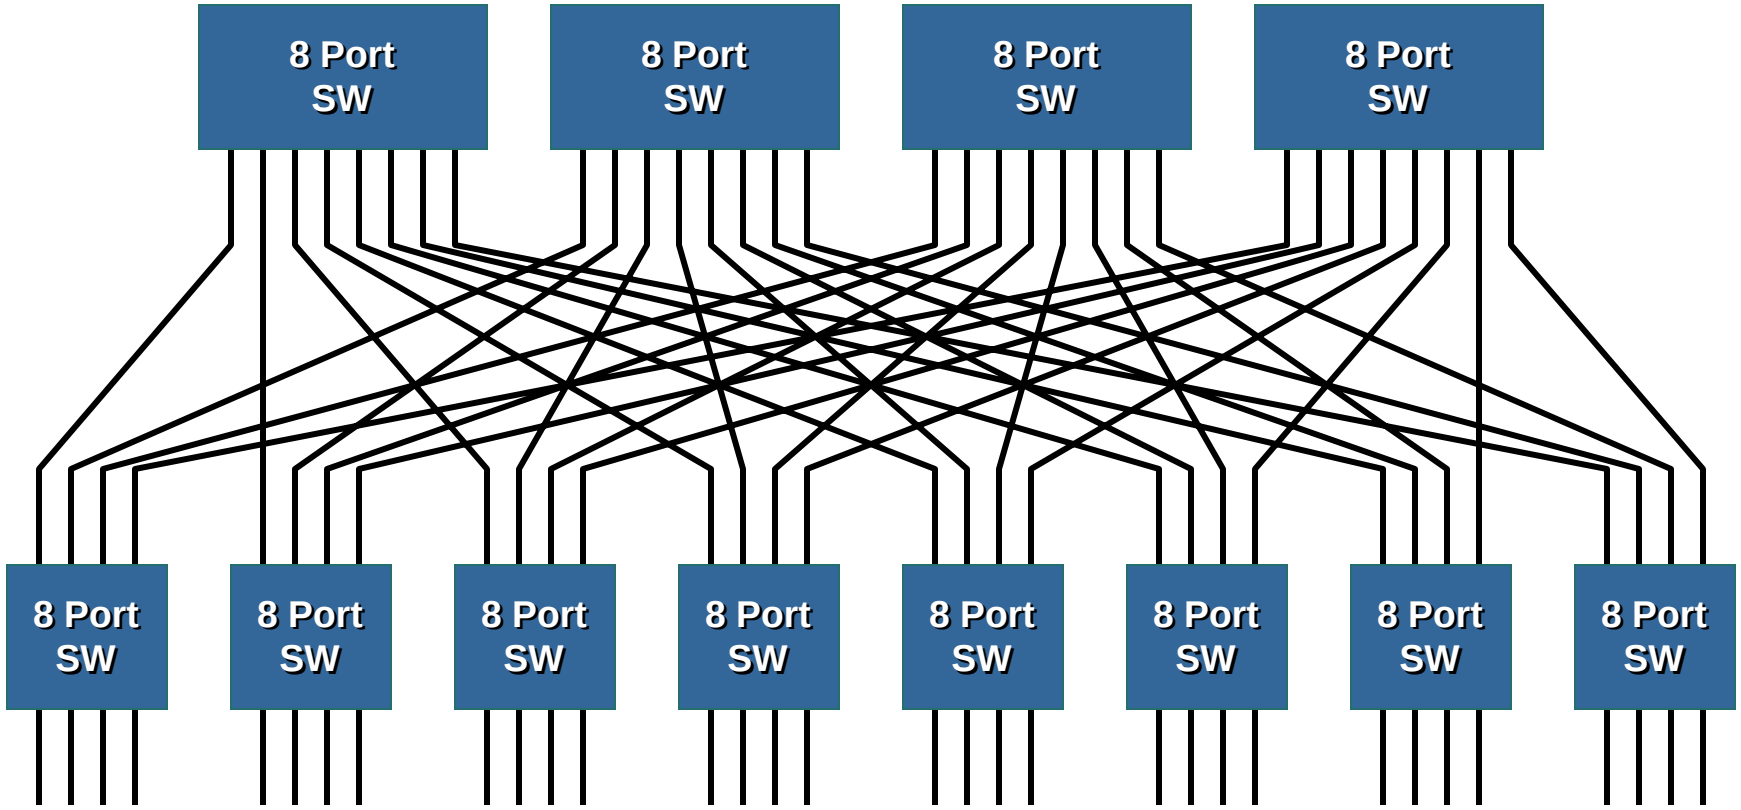

32 Port Switch

Two-Tier

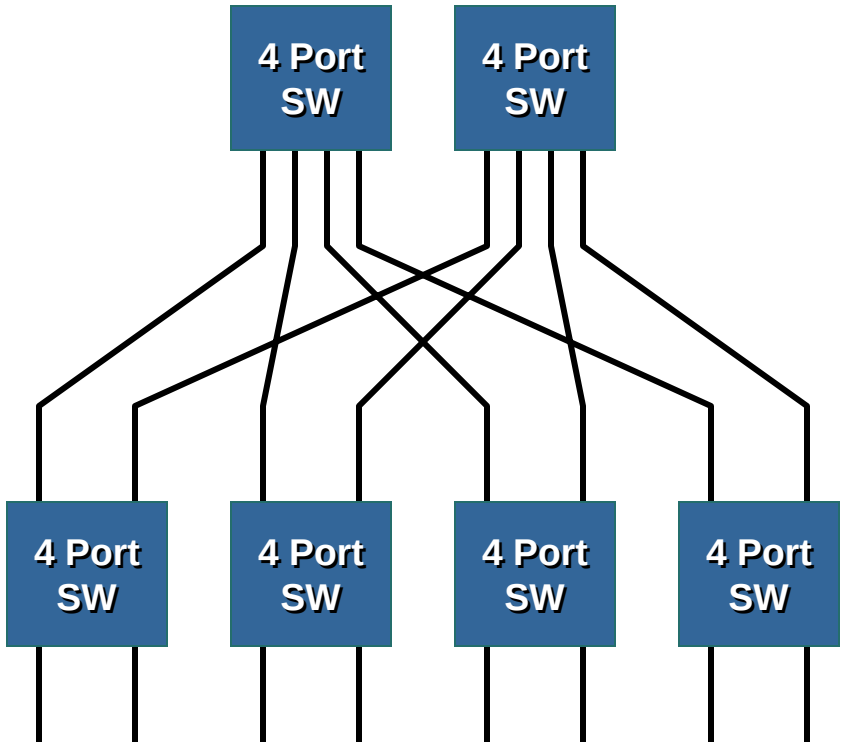

8 Port Switch

Two-Tier

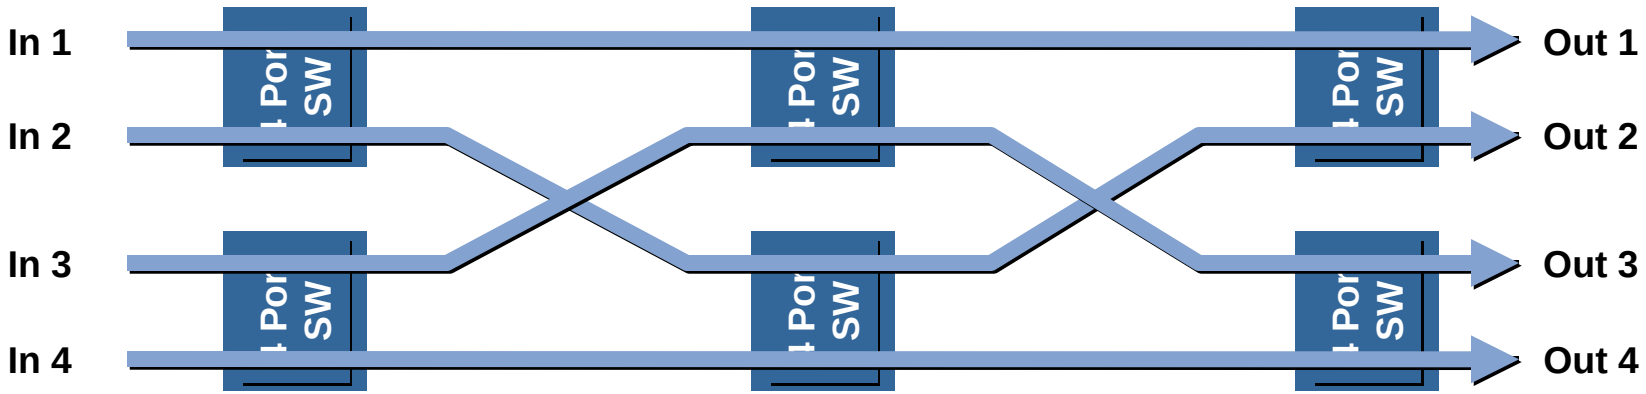

8 Port Switch

Two-Tier

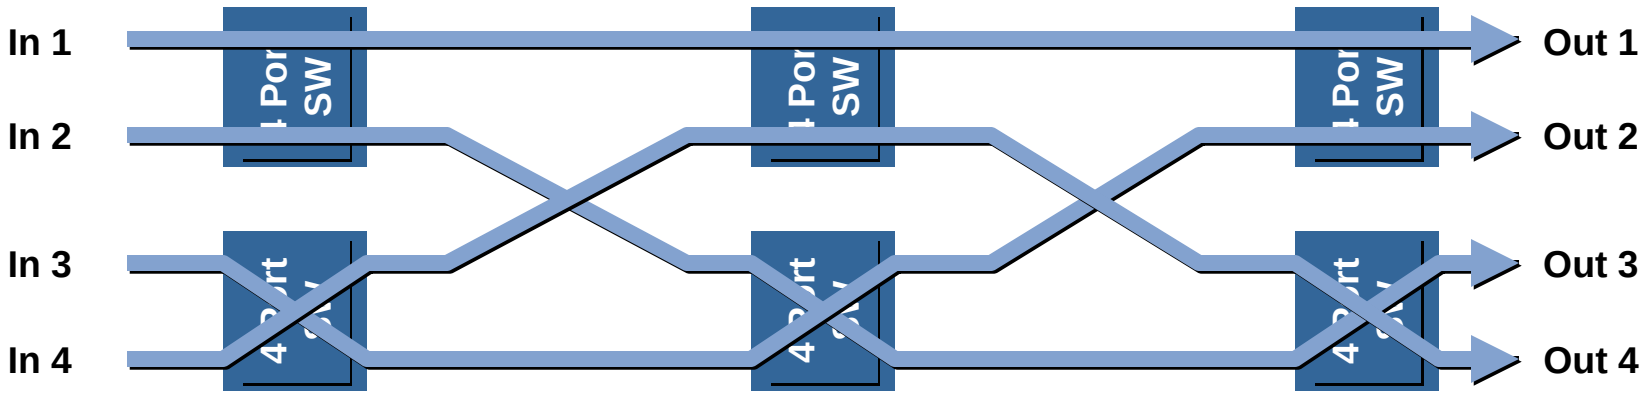

8 Port Switch

Two-Tier

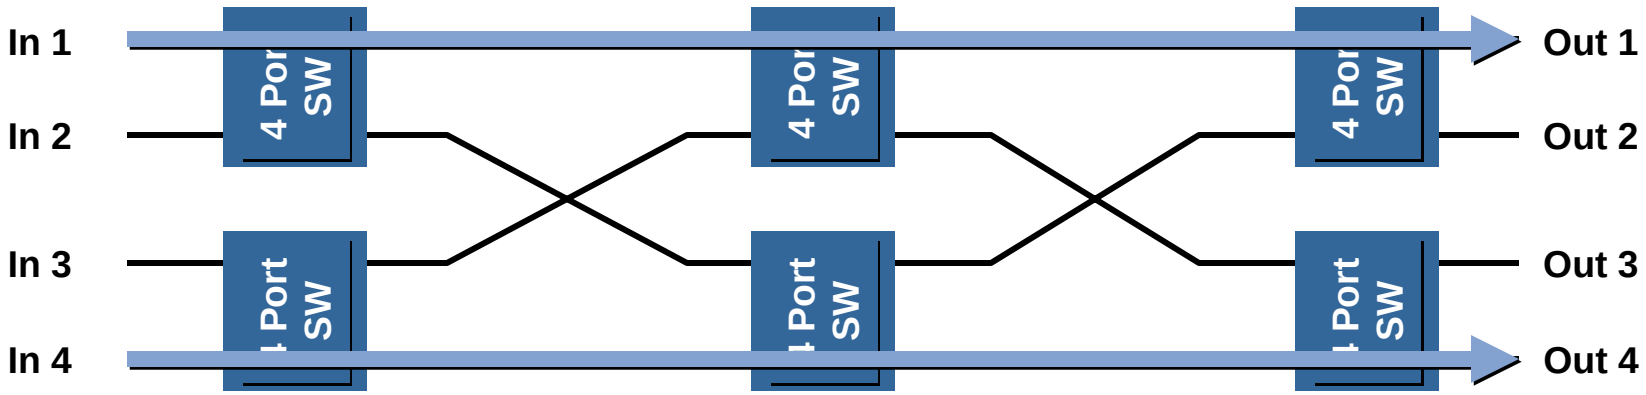

8 Port Switch

Two-Tier

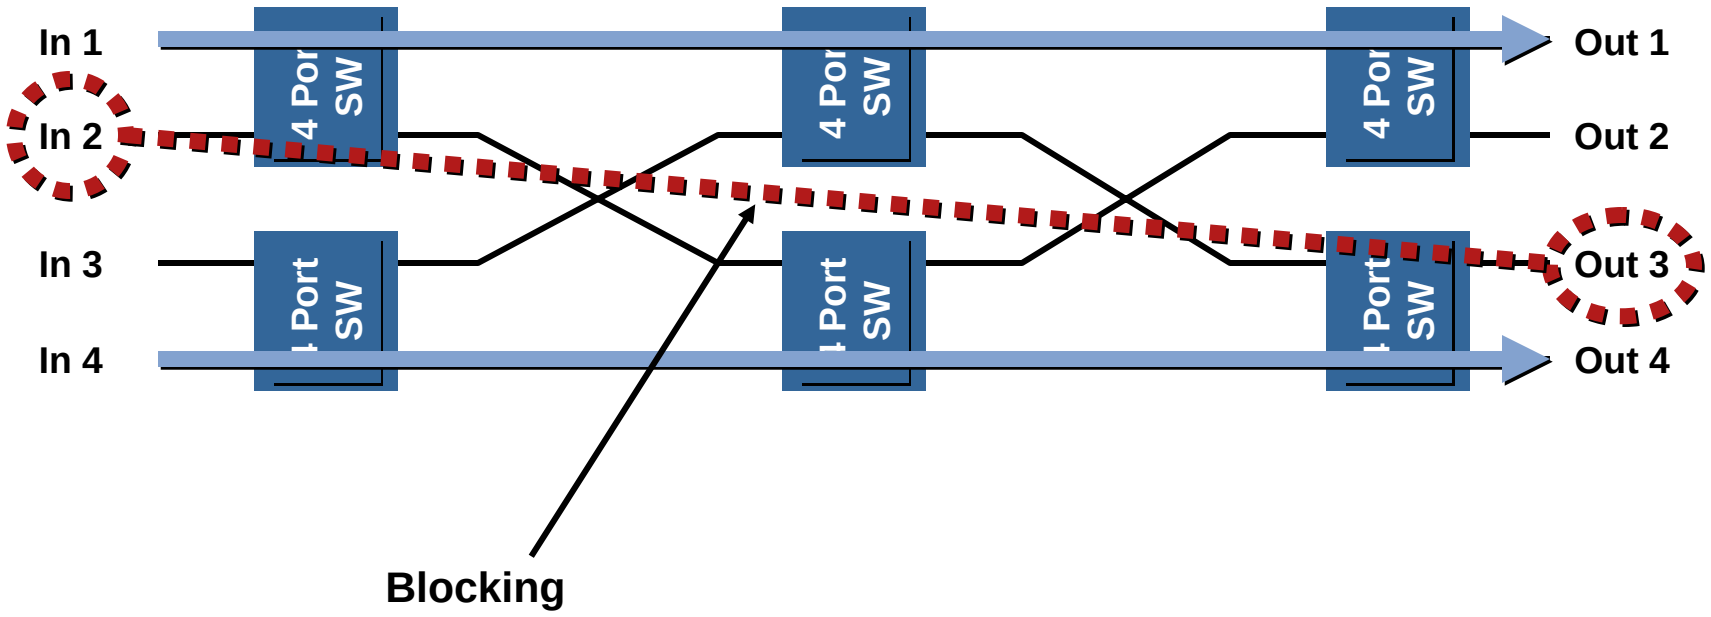

8 Port Switch

Rearrangeably non-blocking

Two-Tier

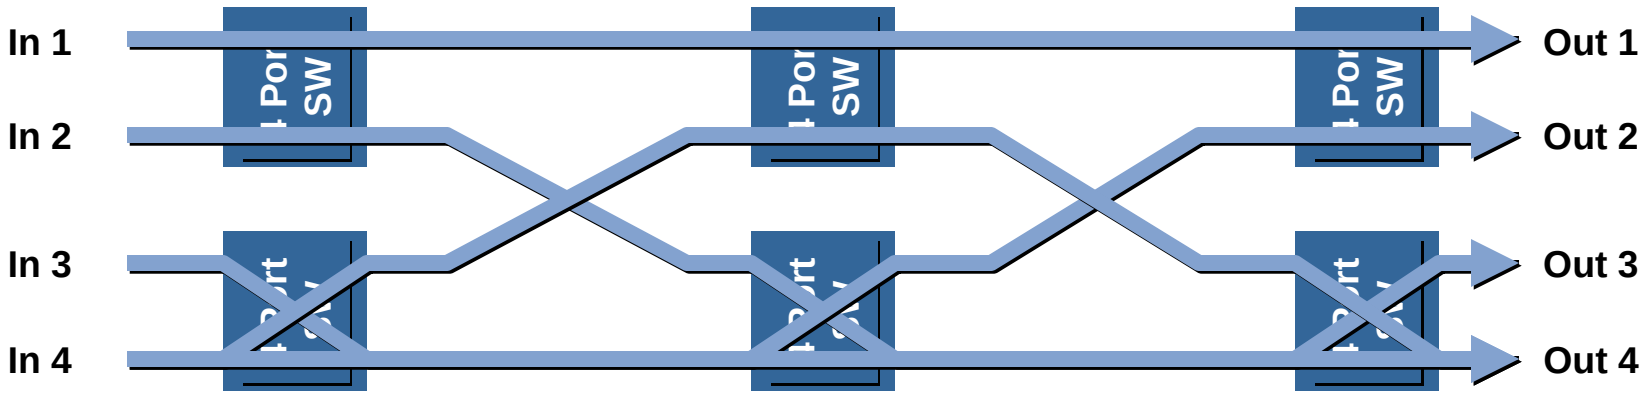

8 Port Switch

Rearrangeably non-blocking

Two-Tier

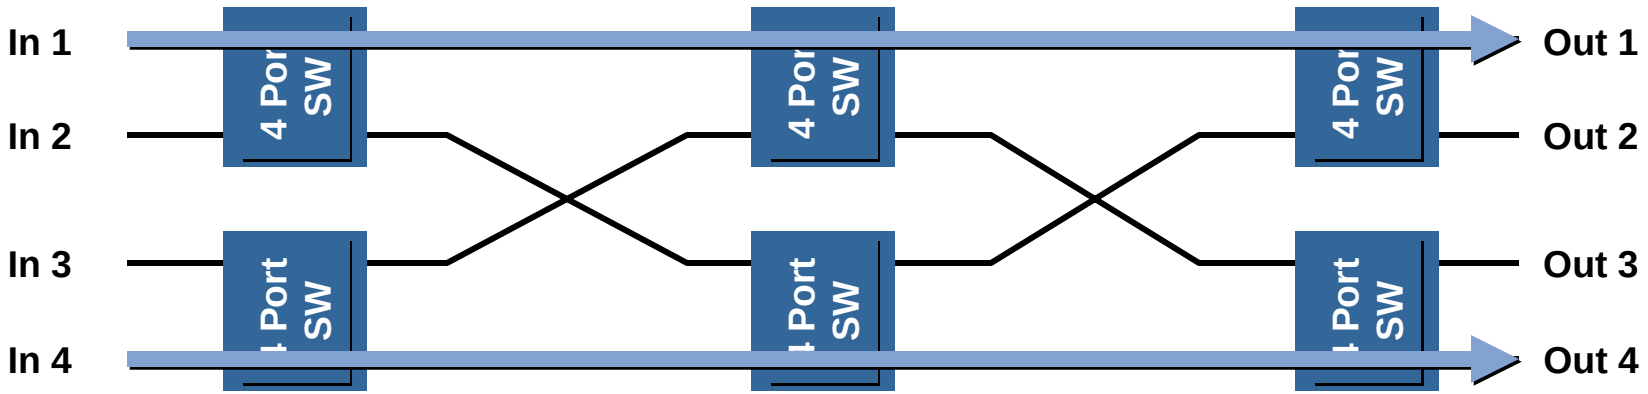

8 Port Switch

Two-Tier

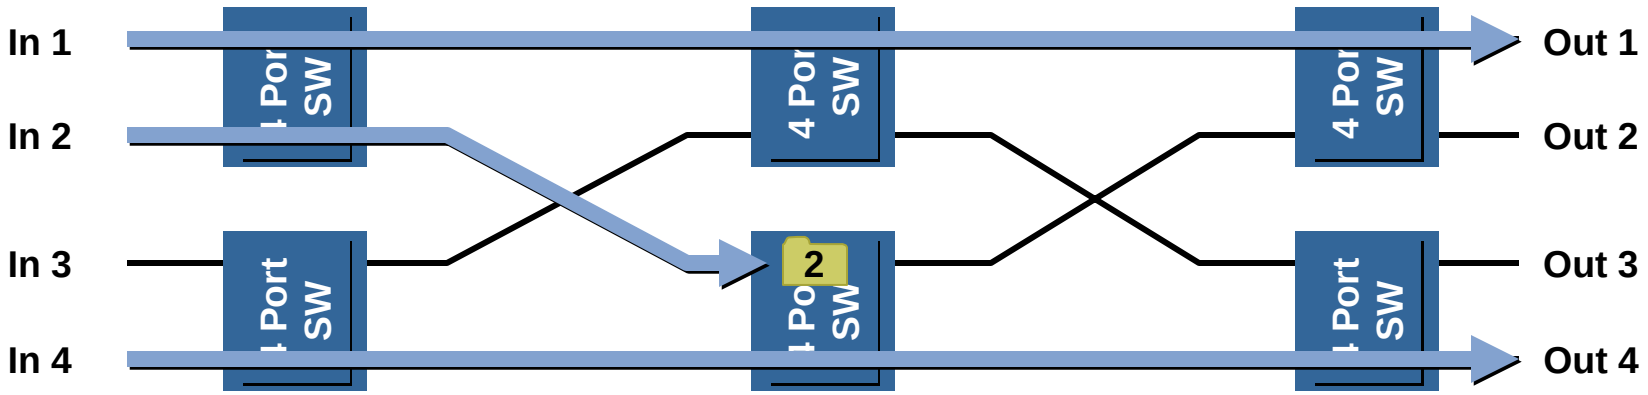

8 Port Switch

Two-Tier

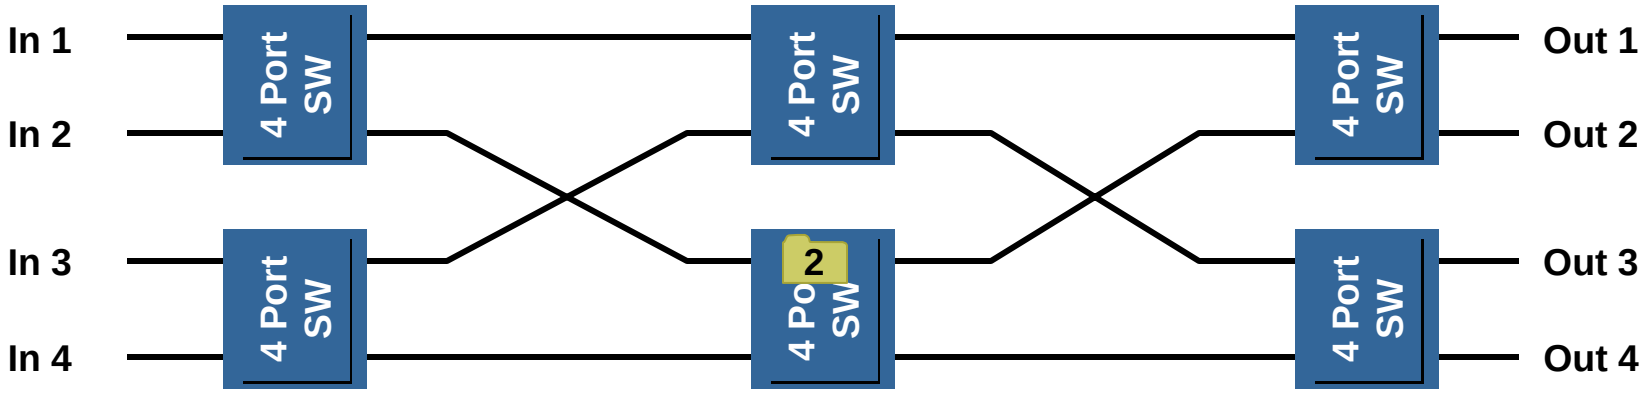

8 Port Switch

Two-Tier

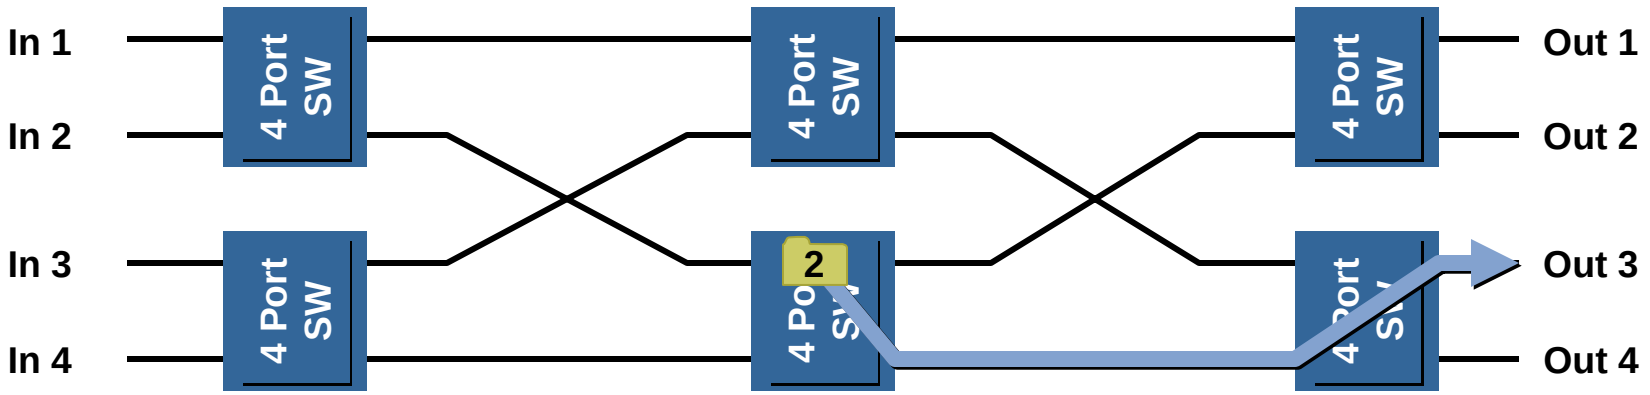

8 Port Switch

Two-Tier-Based Switch—64 Ports

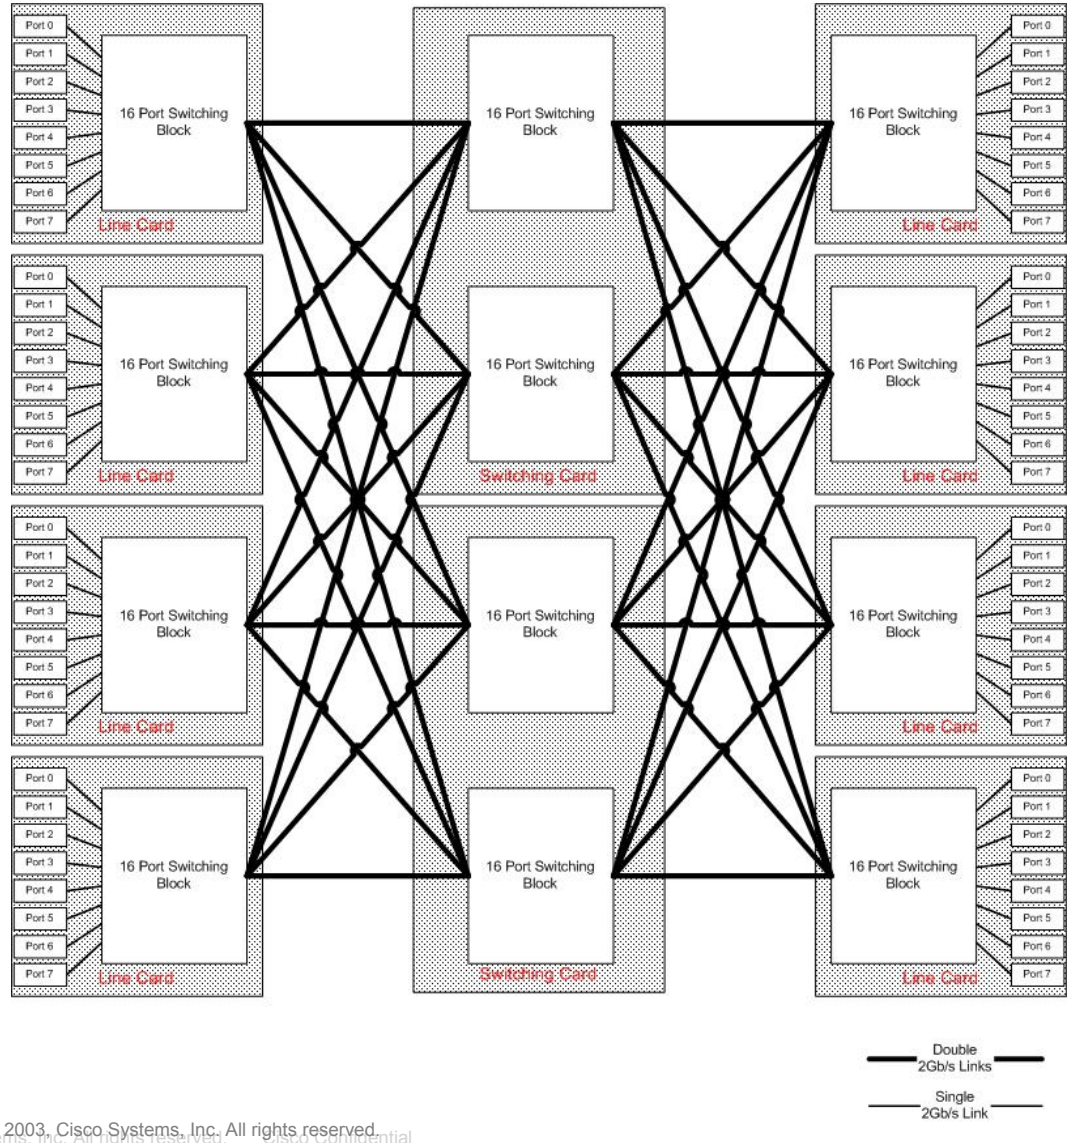

Two-Tier-Based Switch—128 Ports

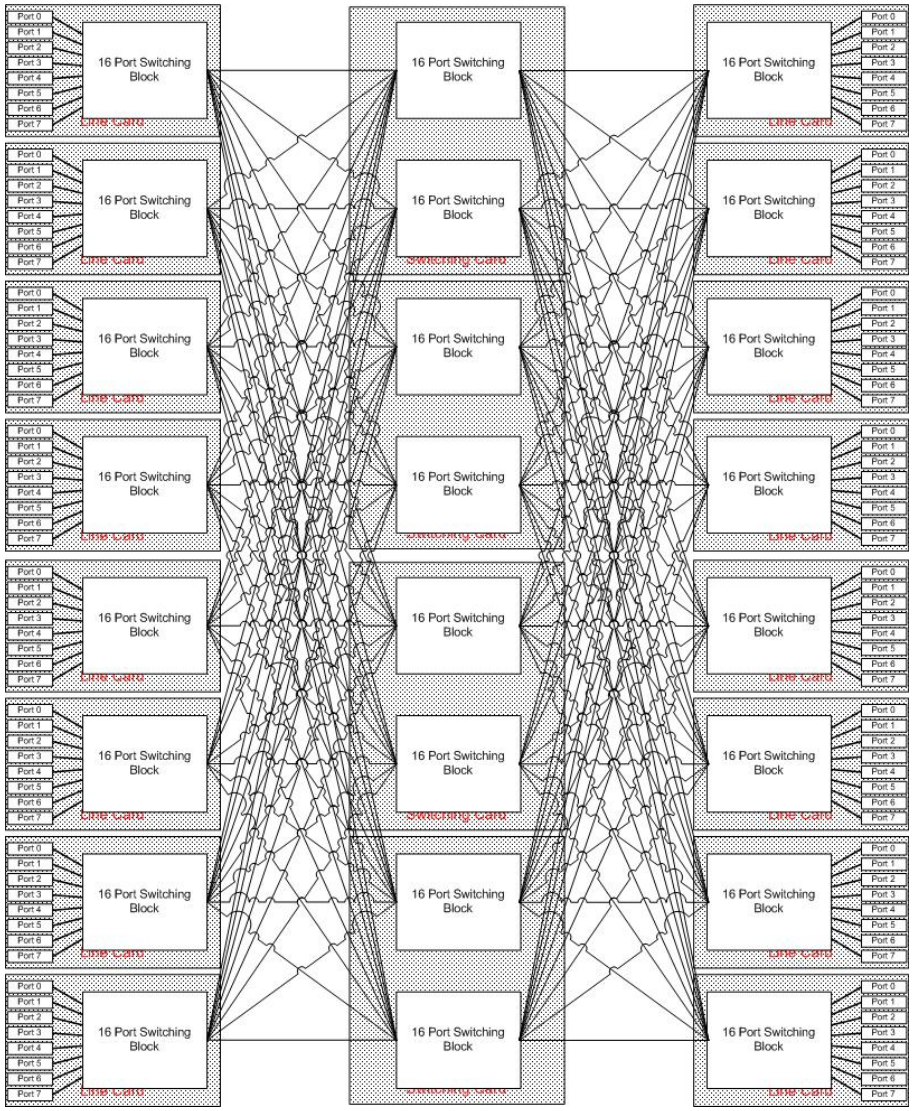

Part II—Architectures

- Centralized Memory
  - Switch on Chip / “System on Chip” (SoC)
- Bus
- Mesh
  - Centralized memory in a mesh
- Two-Tier / Banyan / Clos
  - Centralized memory in two tiers
- **Cross-Bar**

Cross-Bar

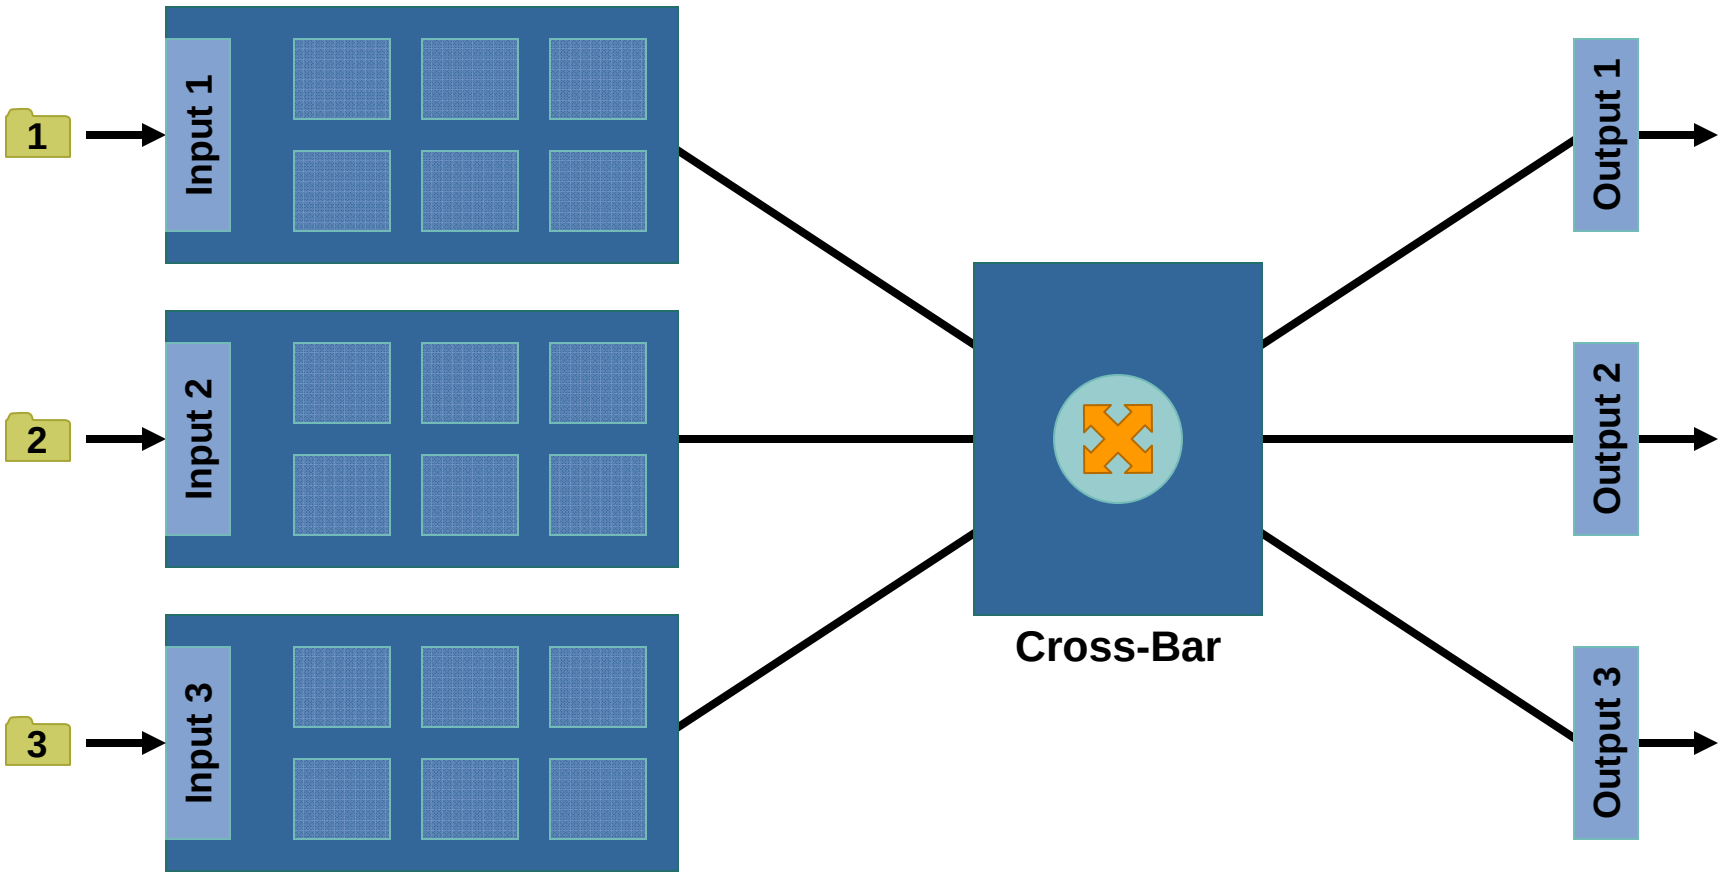

Cross-Bar

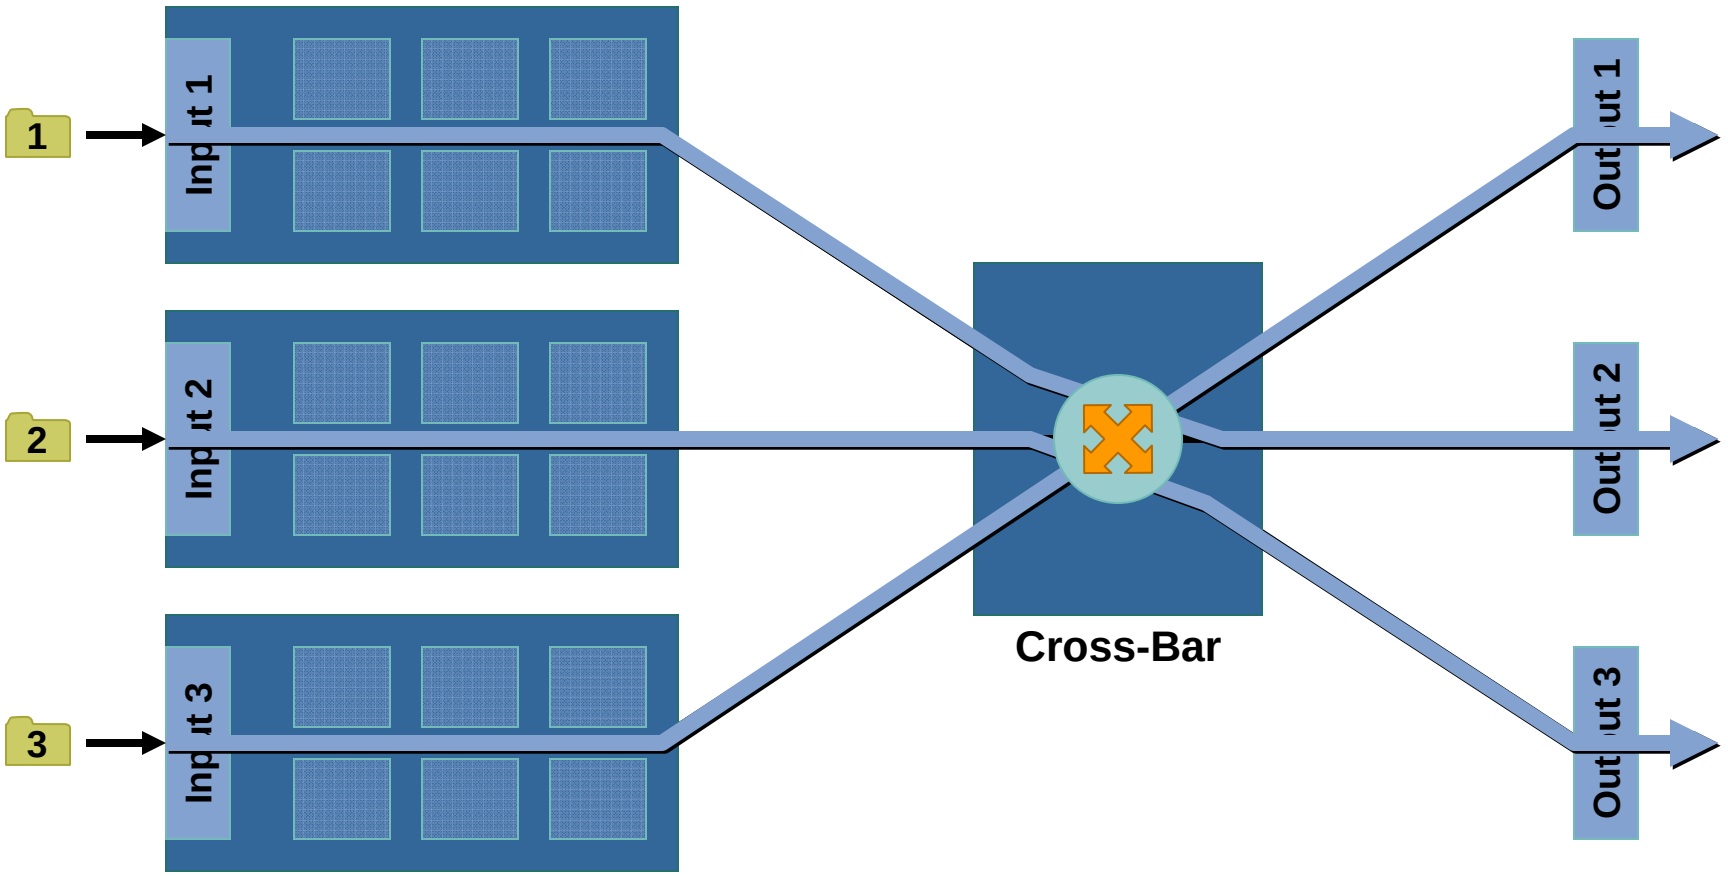

Cross-Bar

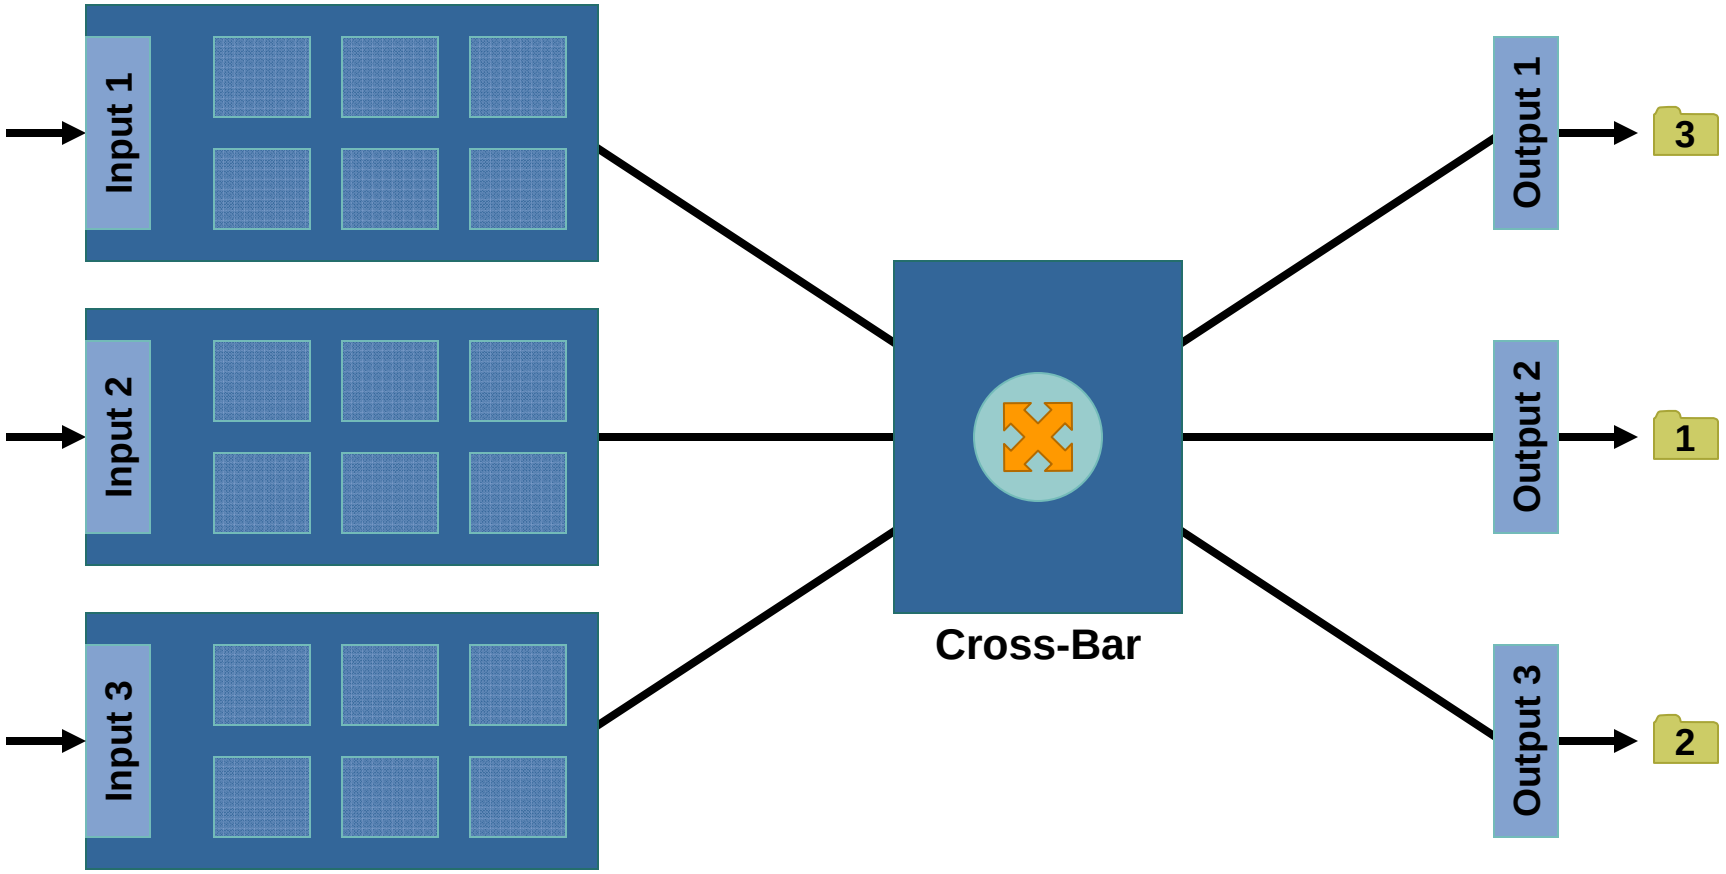

Cross-Bar

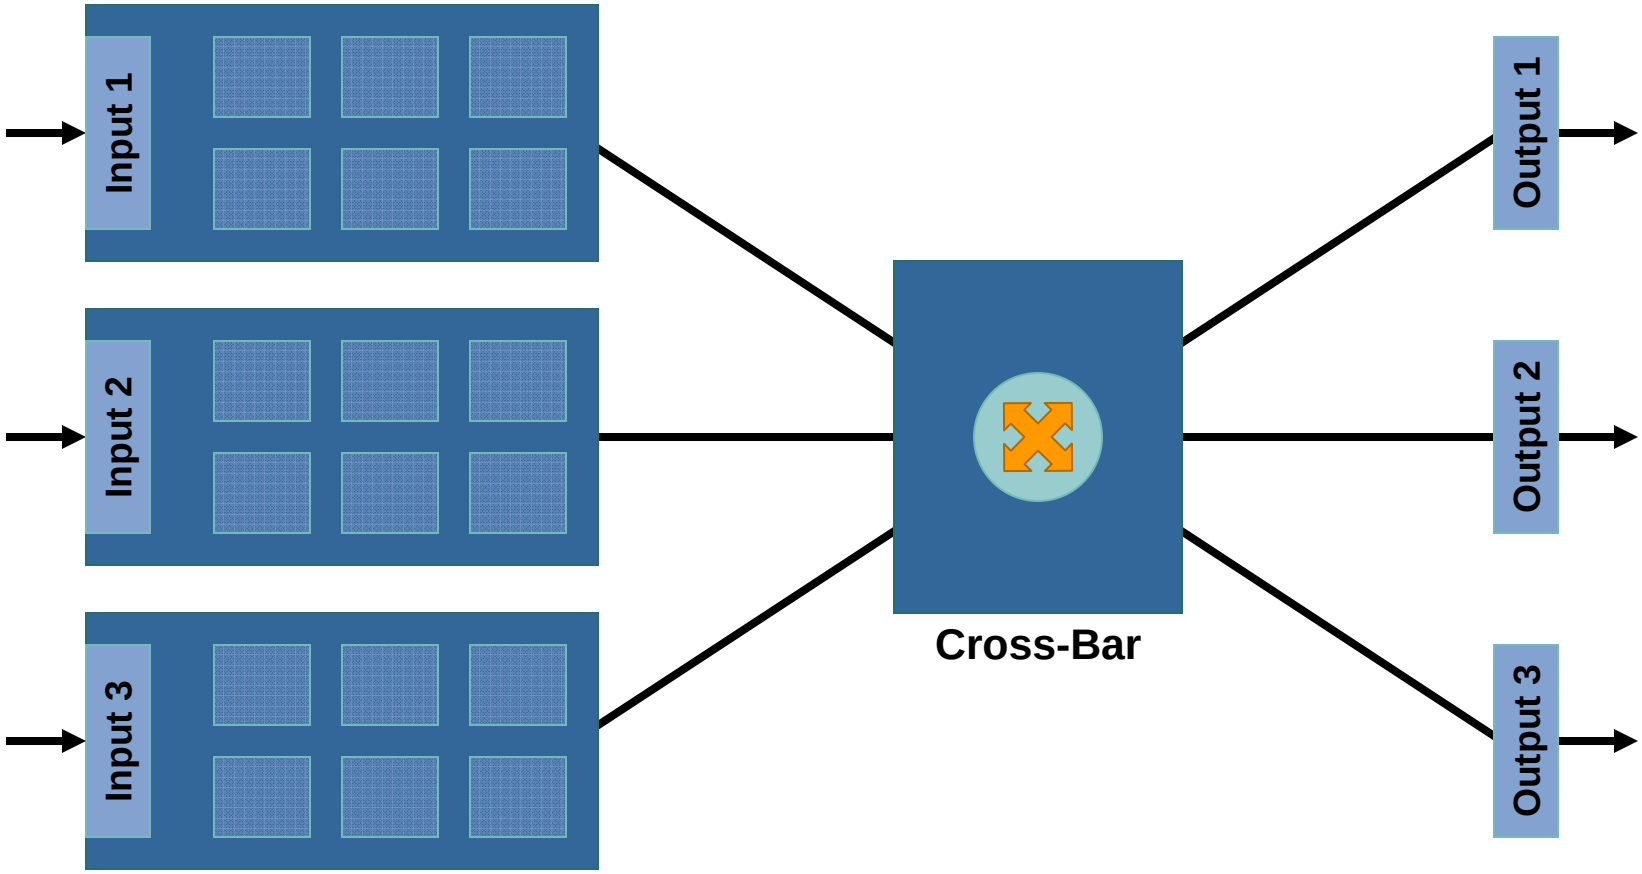

Cross-Bar

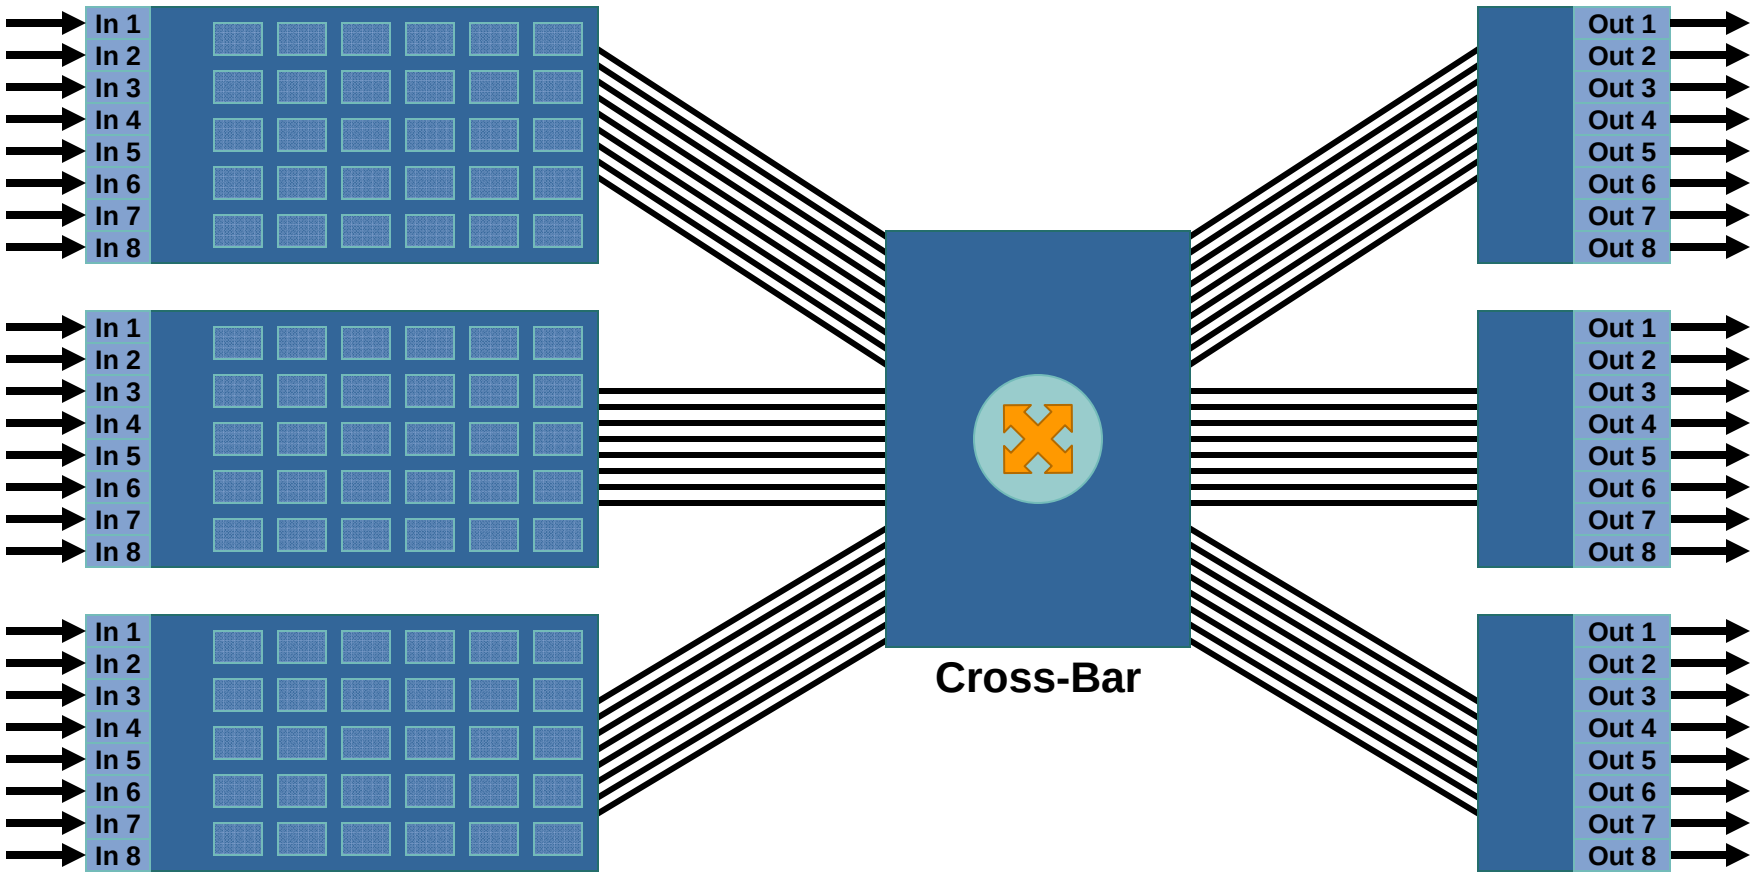

Cross-Bar

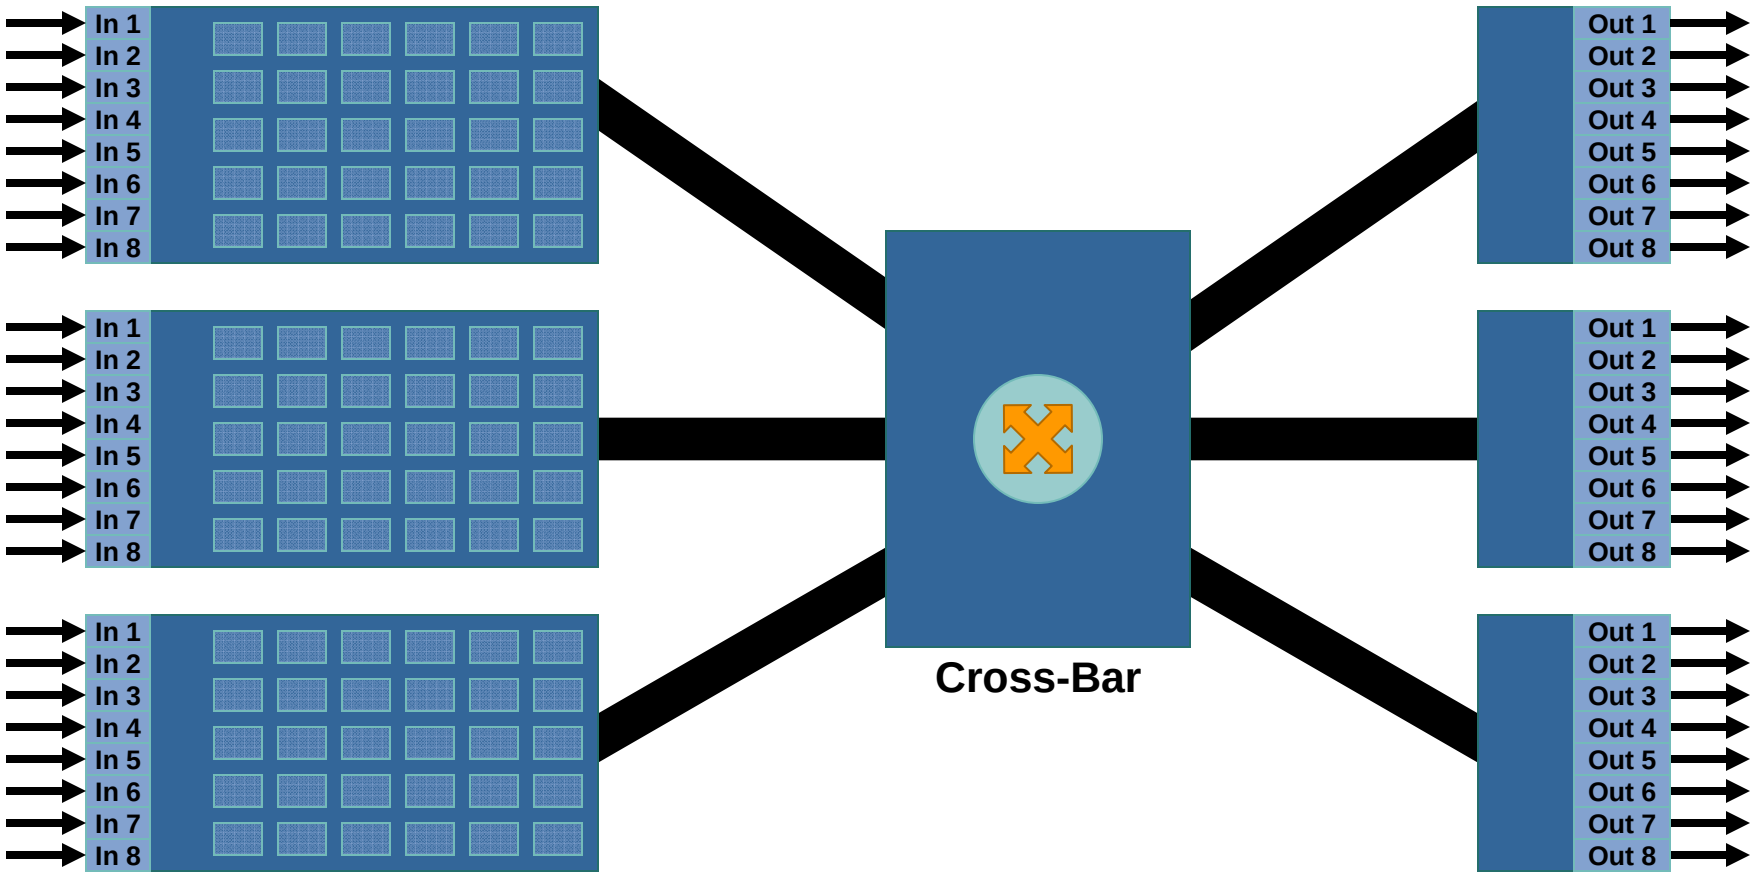

Cross-Bar

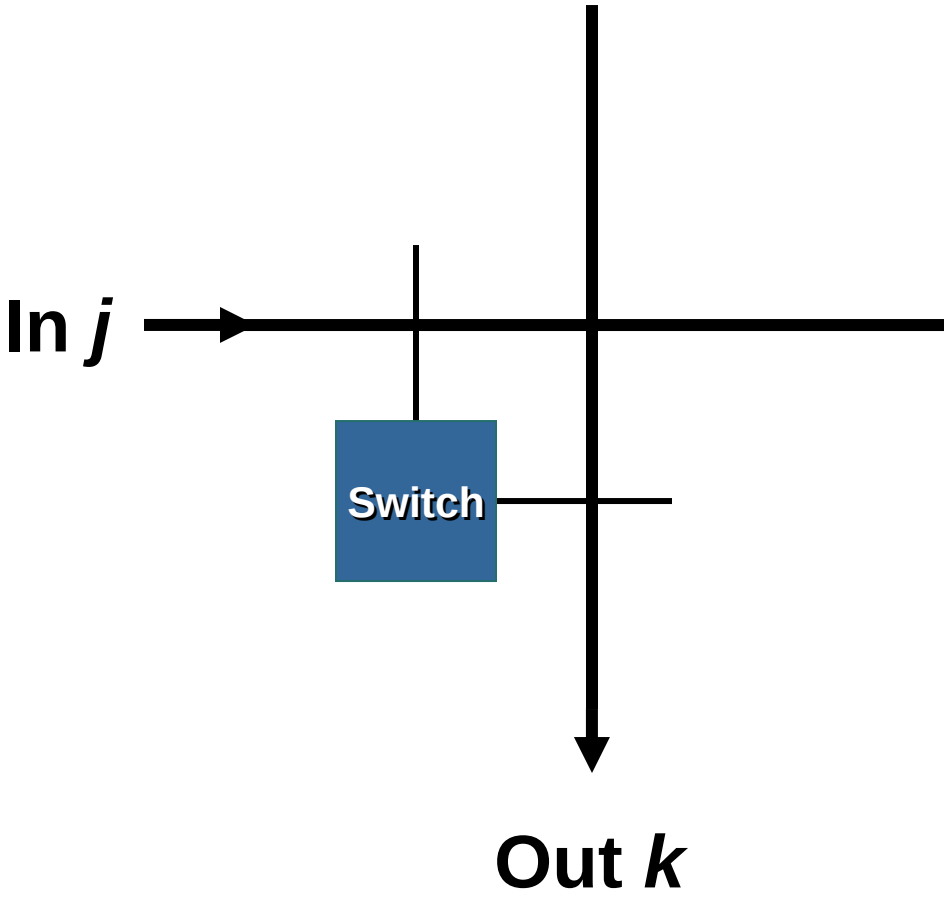

Cross-Bar

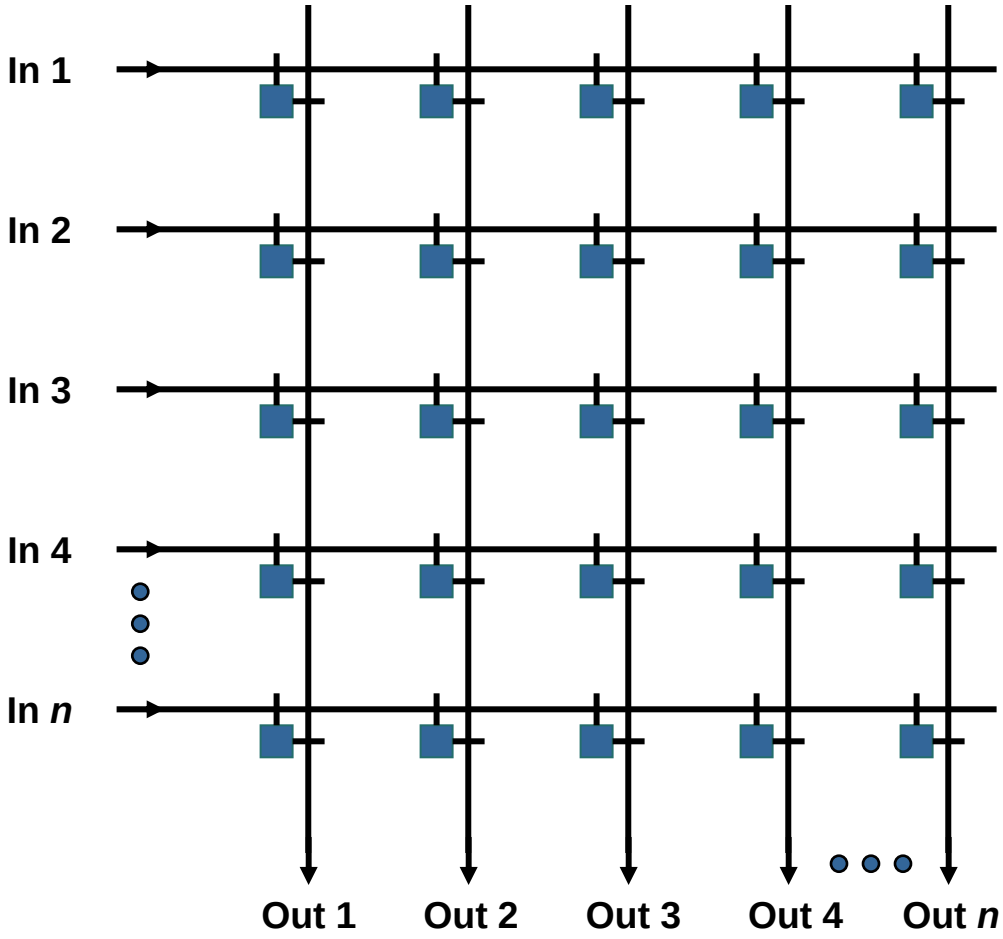

Cross-Bar

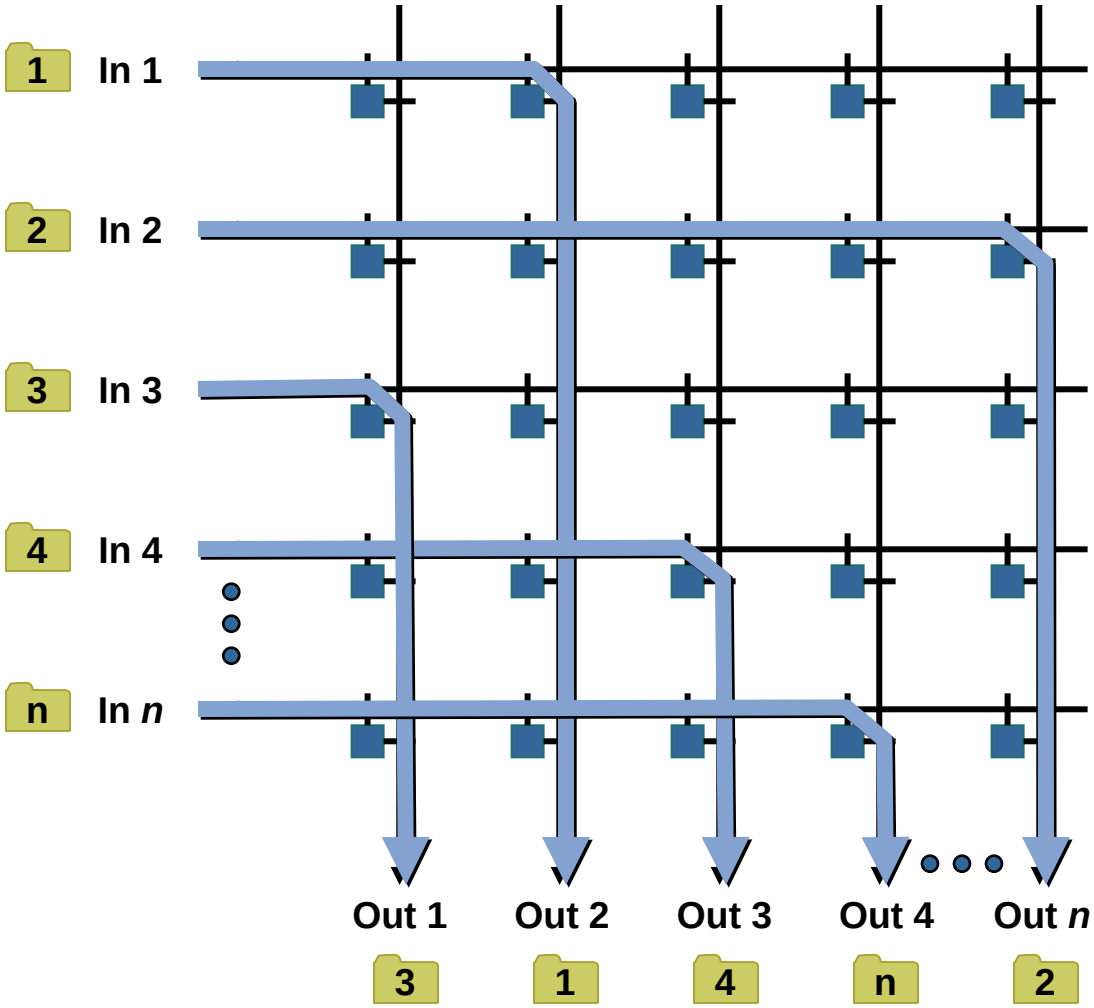

Advantages of Cross-Bar

- Scalable
- Multi-speed capable
- Non-blocking
- Efficient

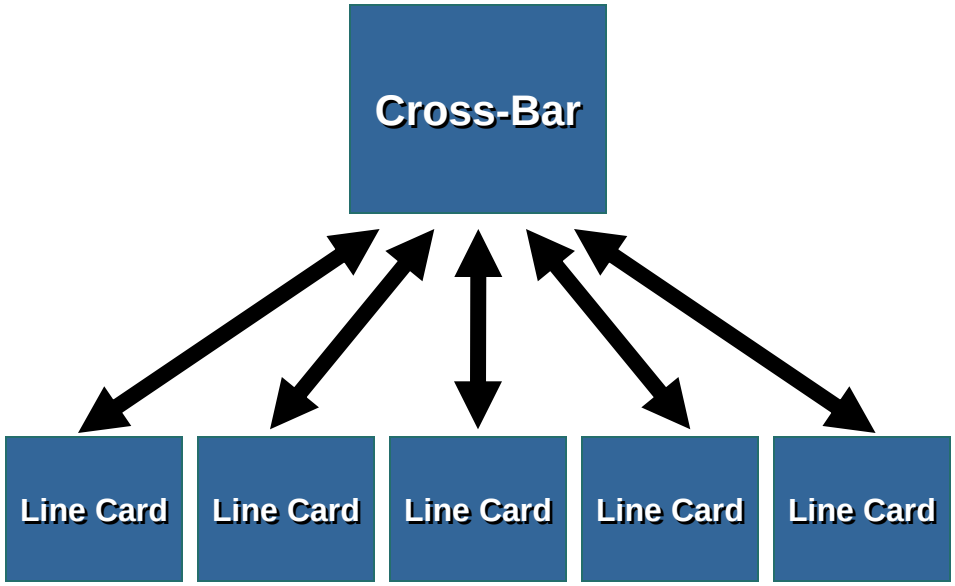

Part III – Real World Implementation

Part III—Real World Implementation

- Cross-bar Performance
  - Classic crossbar vs. Buffered Crossbar
- Virtual Output Queuing
  - Head-of-Line Blocking
- Completing the Switch
- Other Features

Part III—Real World Implementation

- **Cross-bar Performance**
  - Classic crossbar vs. Buffered Crossbar
- Virtual Output Queuing
  - Head-of-Line Blocking
- Completing the Switch
- Other Features

Cross-Bar Model

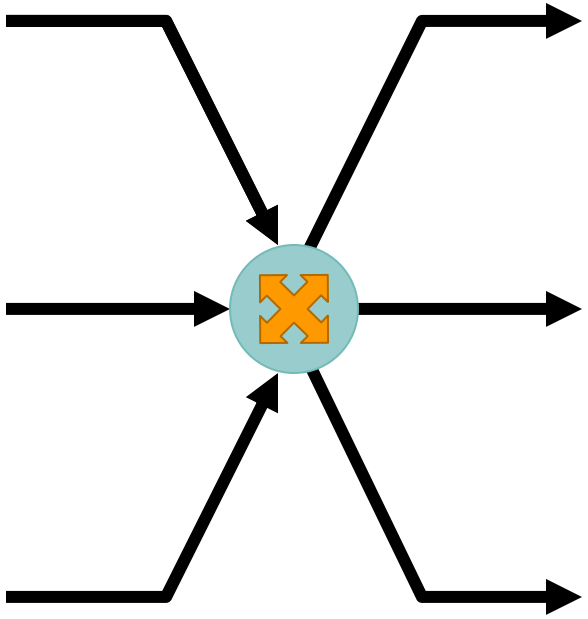

Performance Issue with Cross-Bars

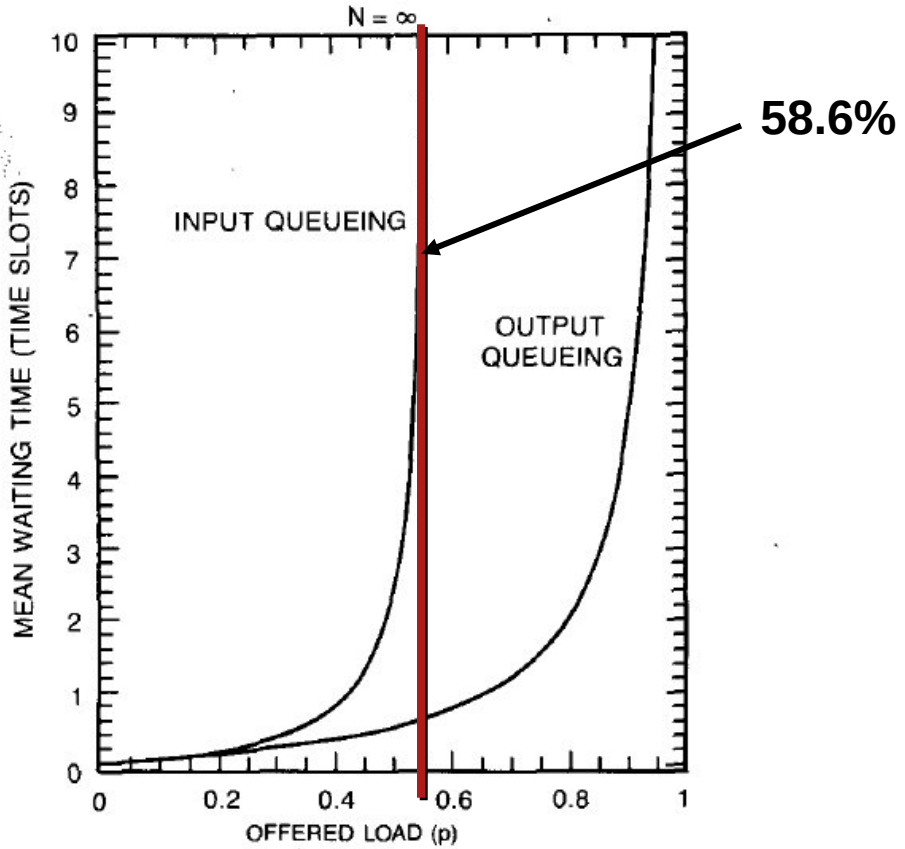

Fig. 9. A comparison of the mean waiting time for input queueing and output queueing for the limiting case of $N = \infty$.

Source: M. J. Karol, M.G. Hluchyj, S. P. Morgan, "Input Versus Output Queueing [sic] on a Space-Division Packet Switch", IEEE Transactions on Communications, Vol COM-35, No 12, December 1987, page 1353

Cat6K & MDS Family Cross-Bar

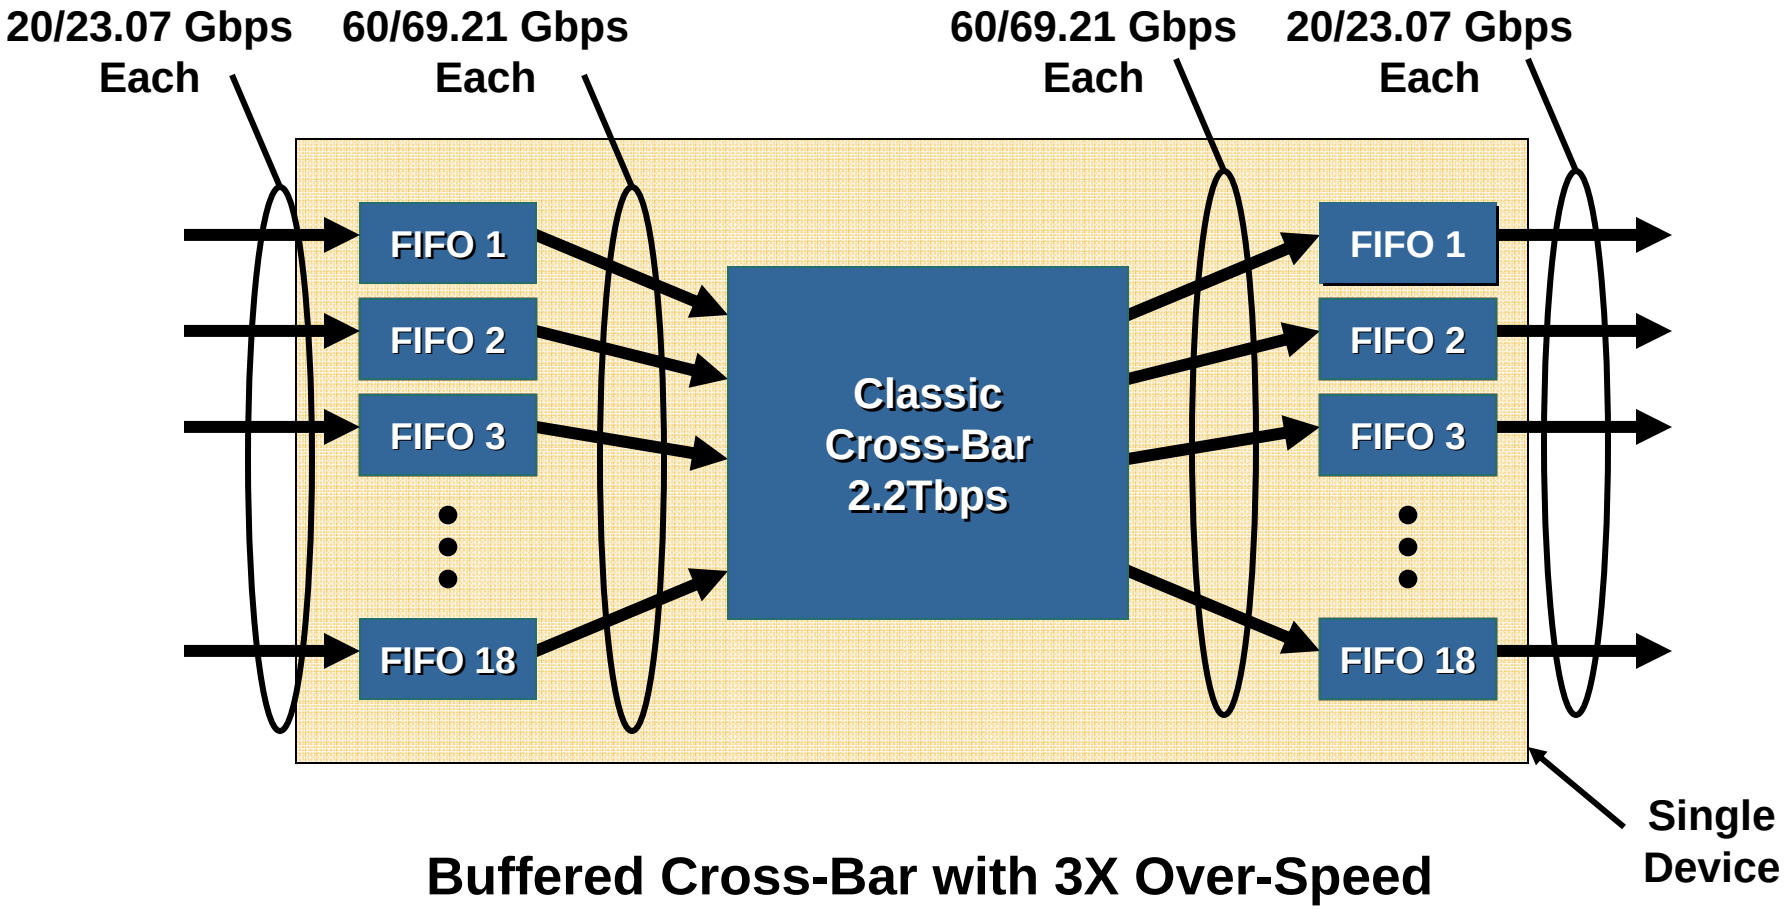

Buffered Cross-Bar with 3X Over-Speed

Part III—Real World Implementation

- Cross-bar Performance
  - Classic crossbar vs. Buffered Crossbar
- **Virtual Output Queuing**
  - Head-of-Line Blocking
- Completing the Switch
- Other Features

Buffering—Input Buffering

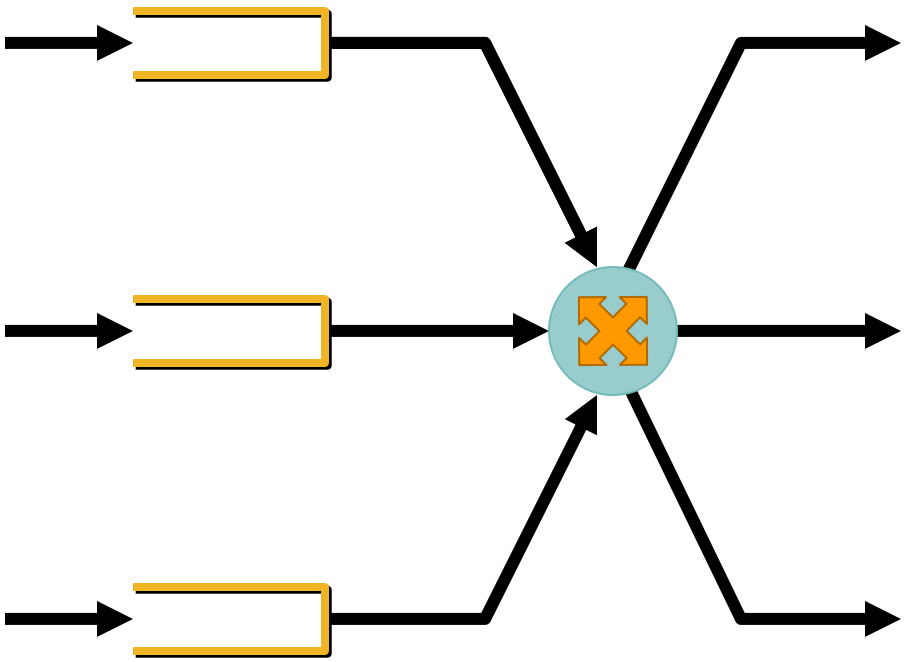

Head of Line Blocking

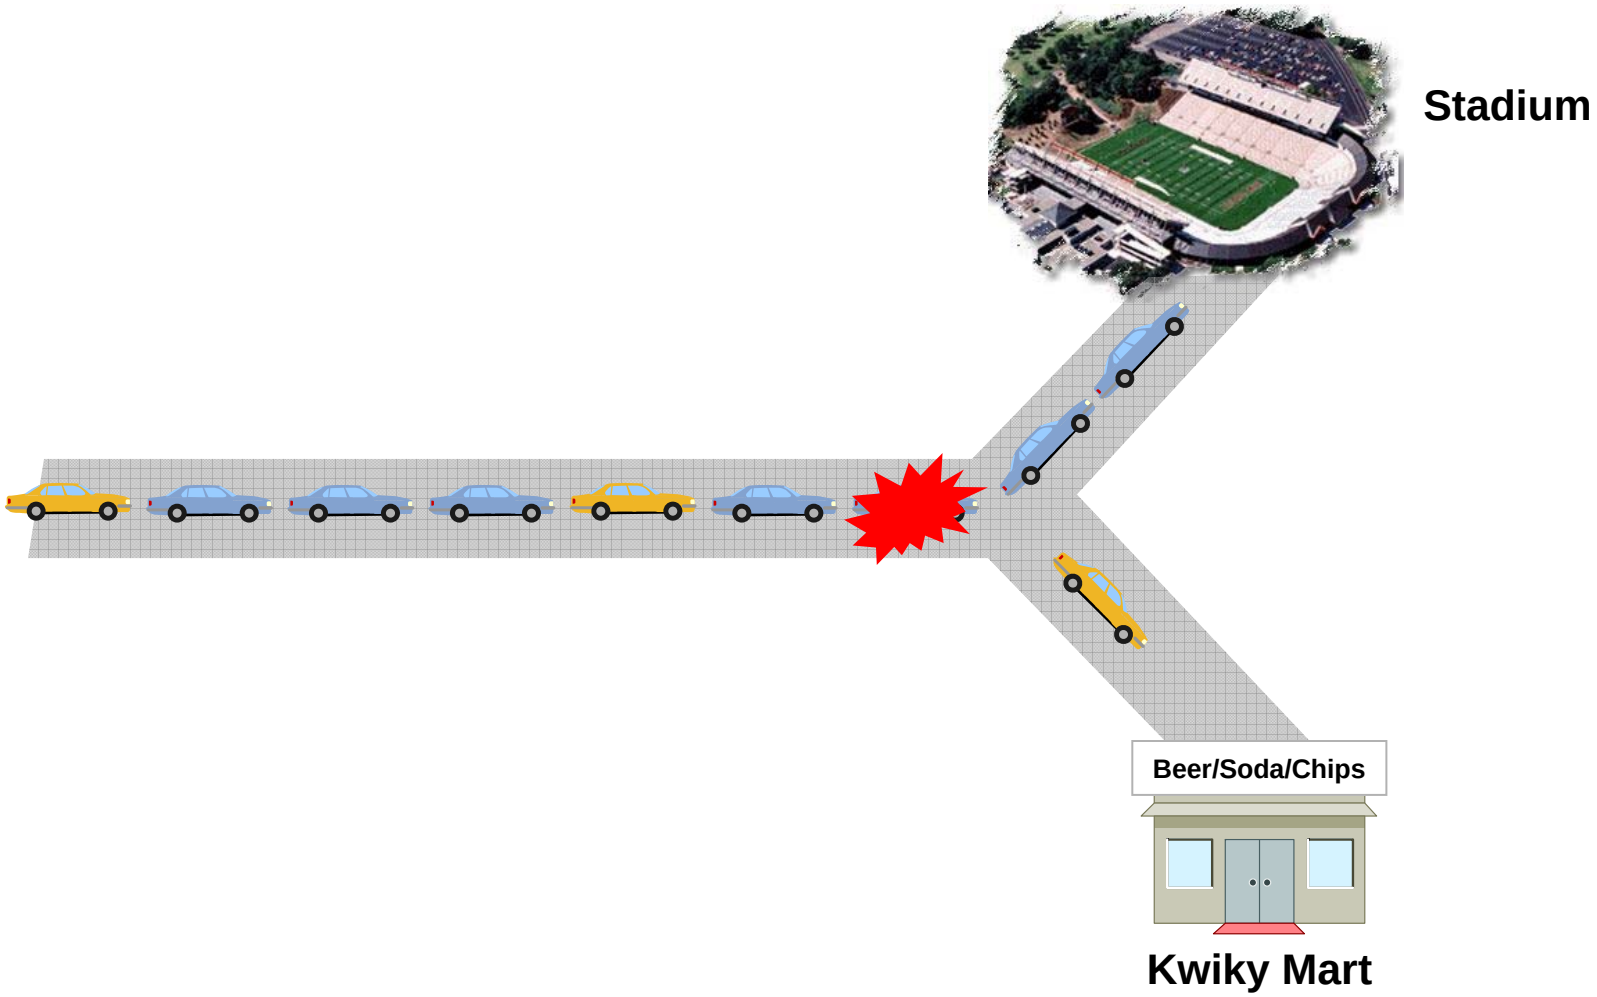

Head of Line Blocking

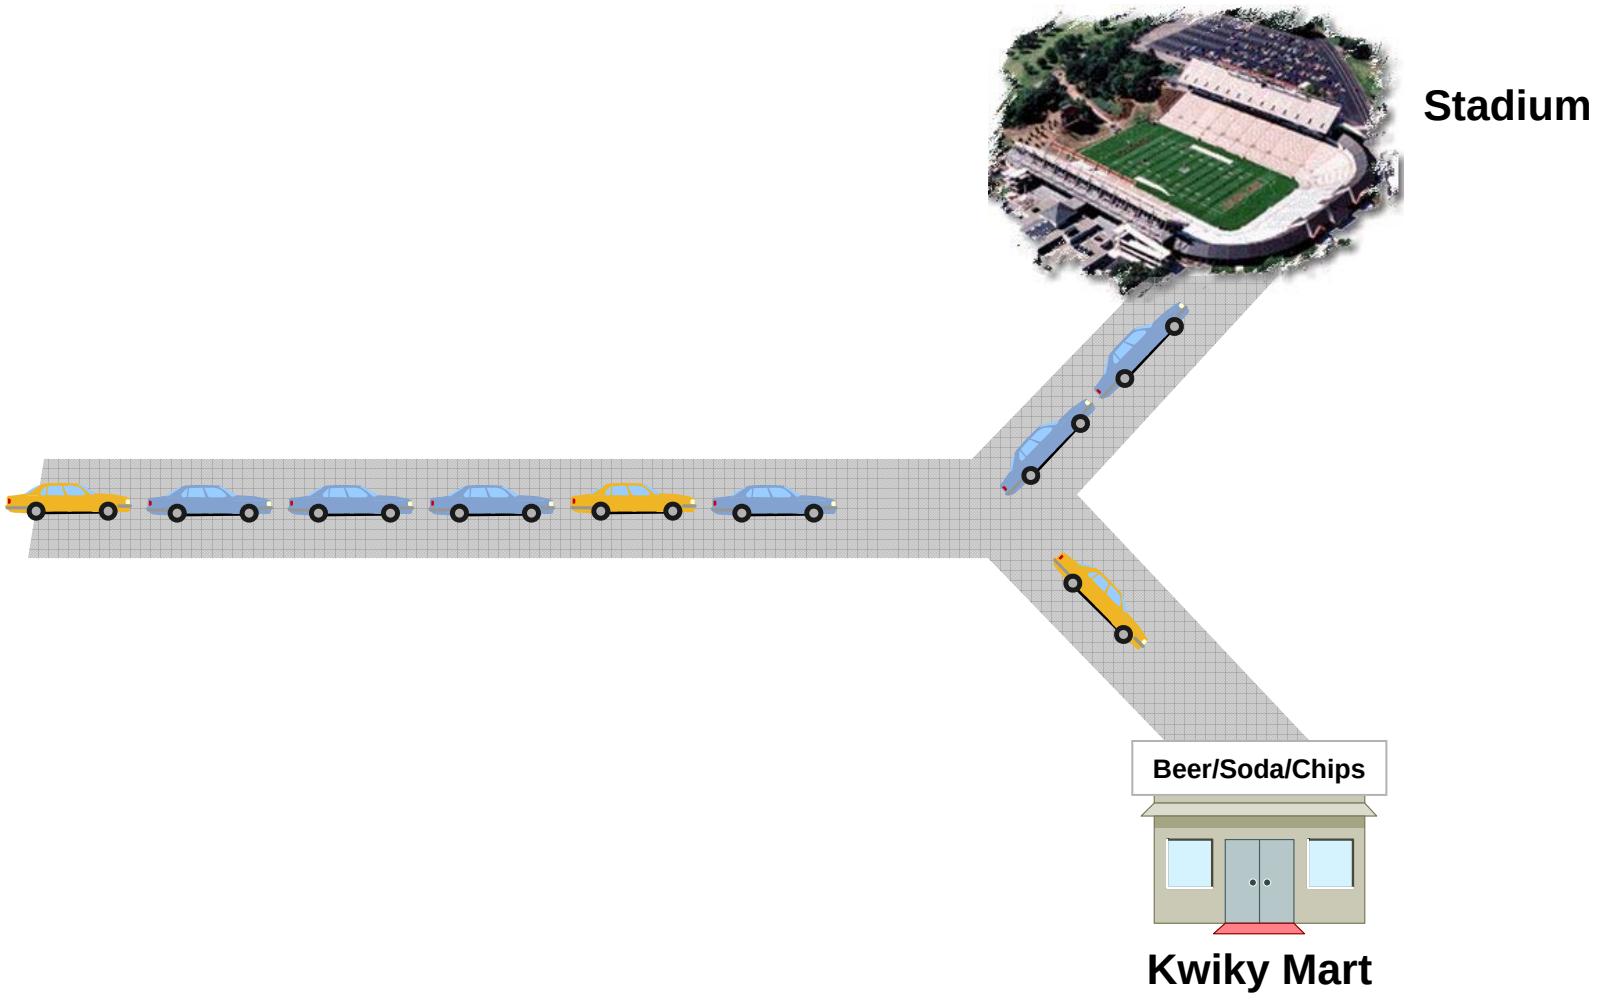

Output Queuing

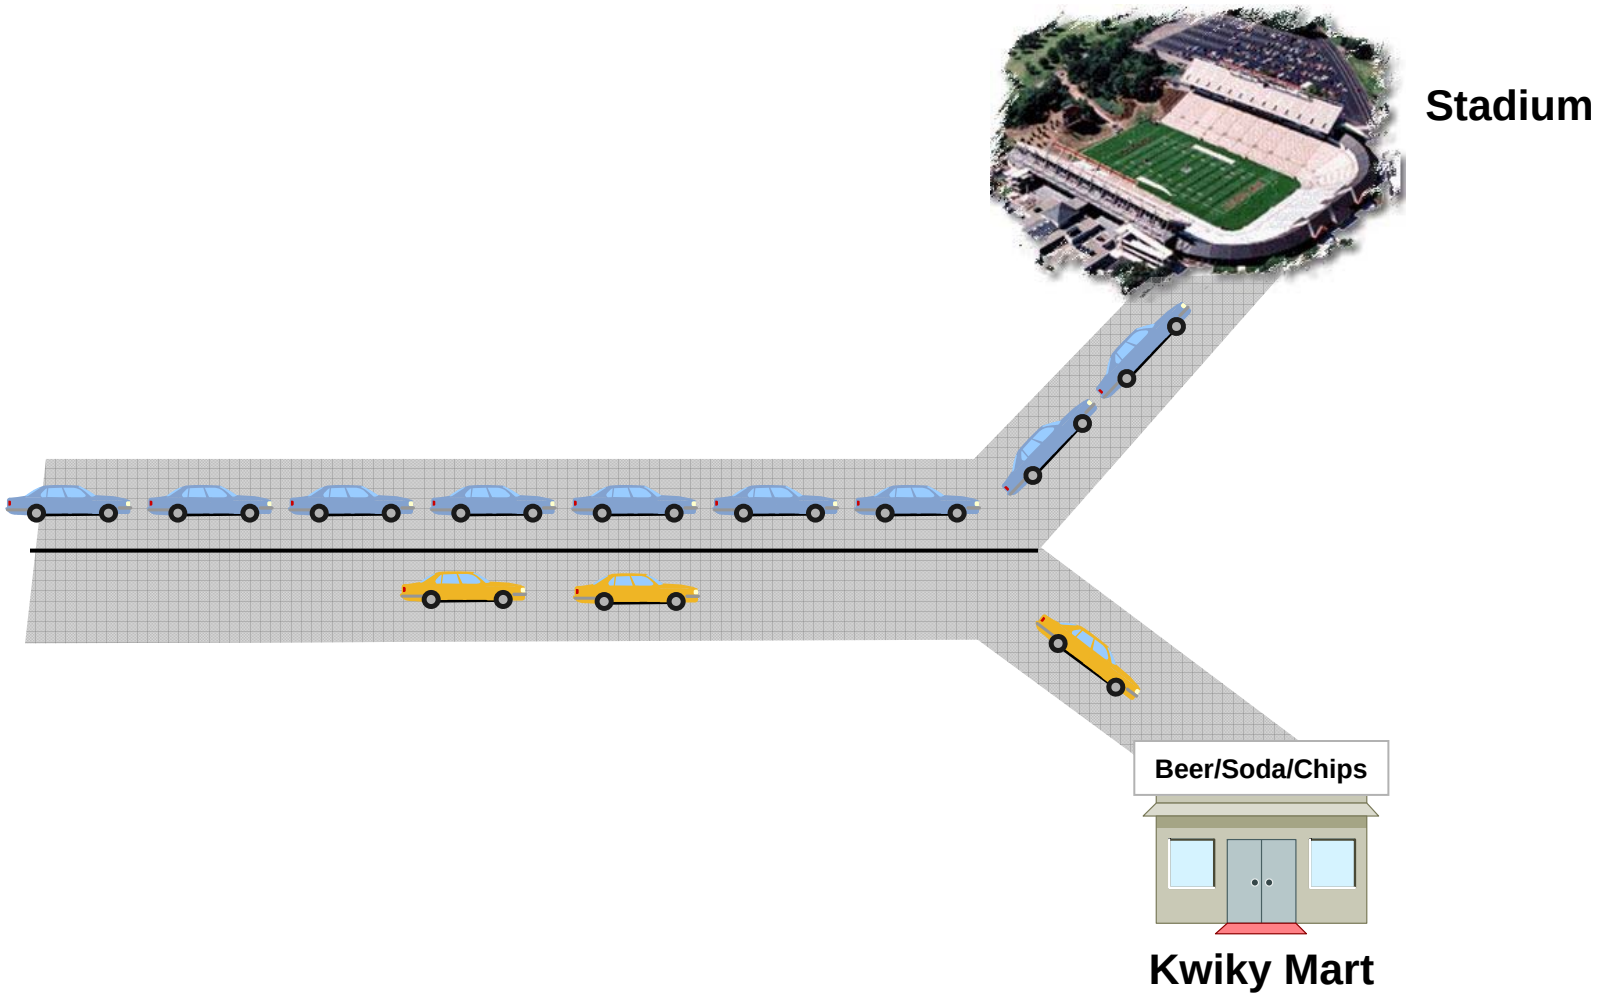

Head of Line Blocking

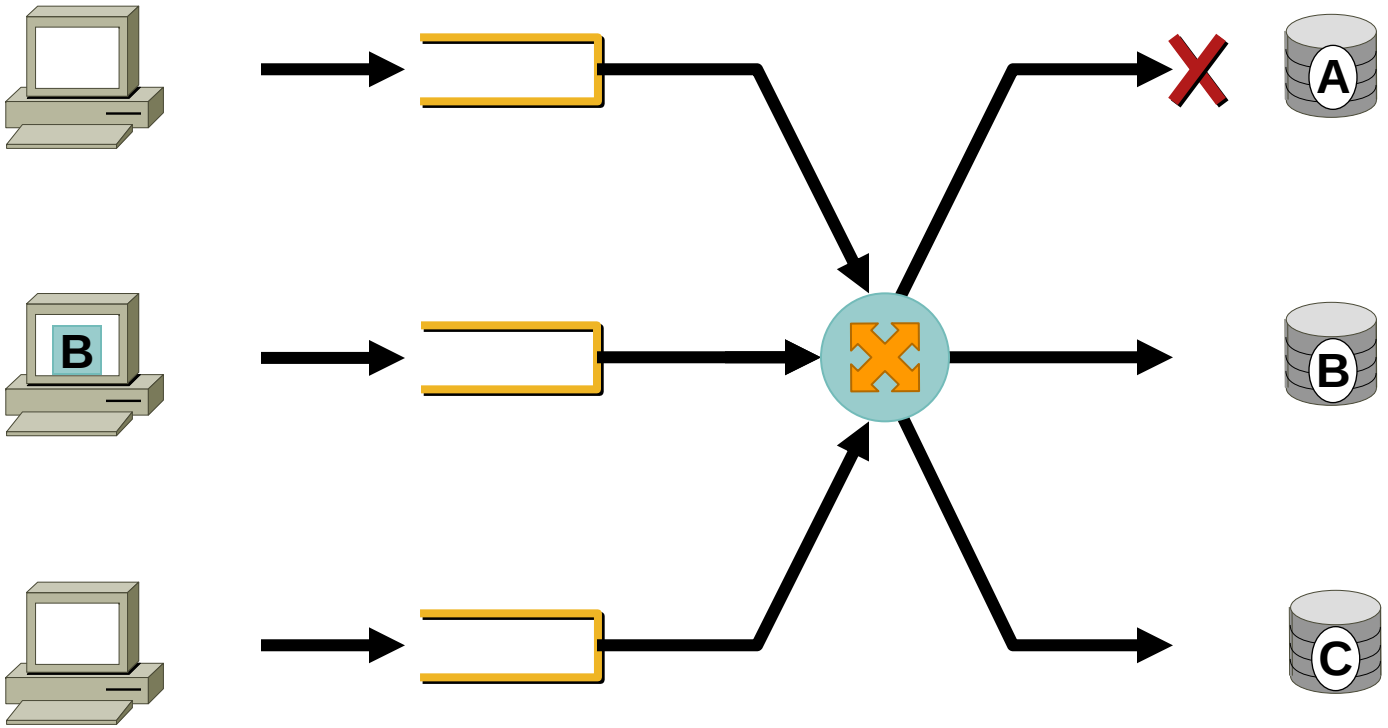

Head of Line Blocking

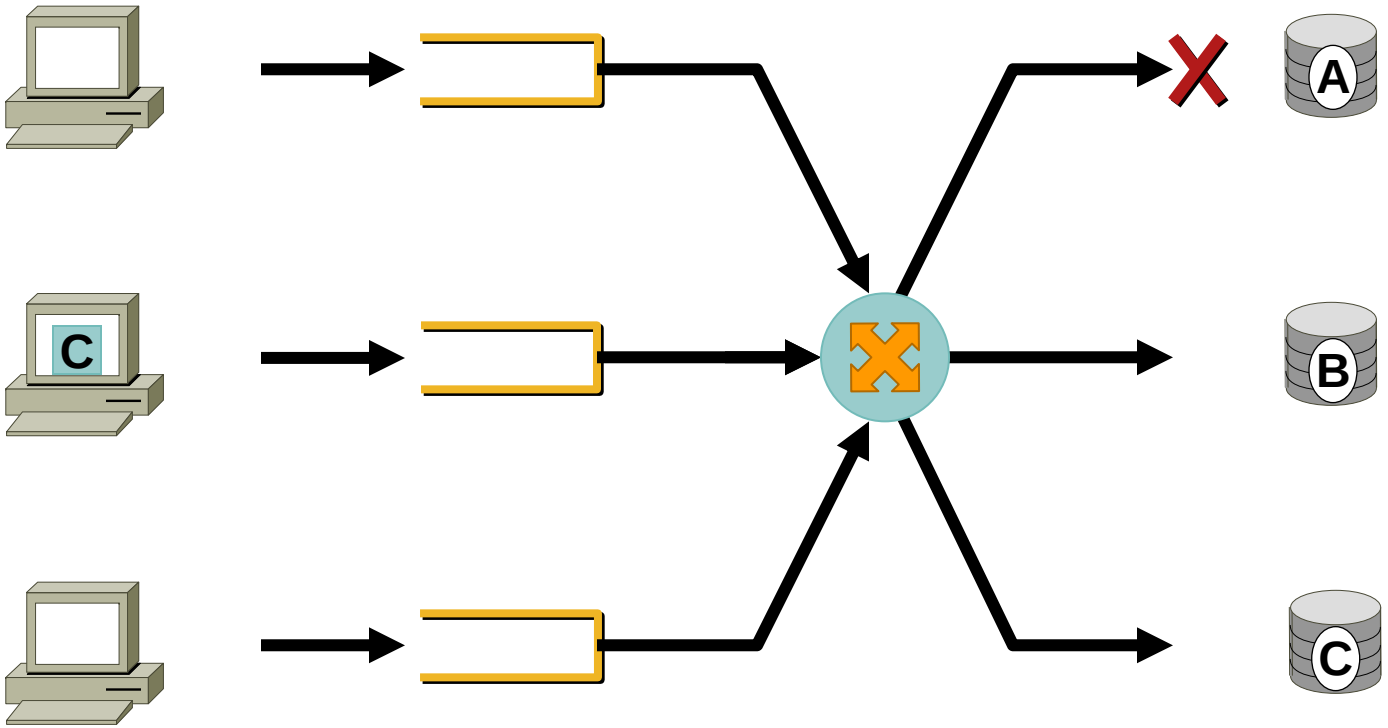

Head of Line Blocking

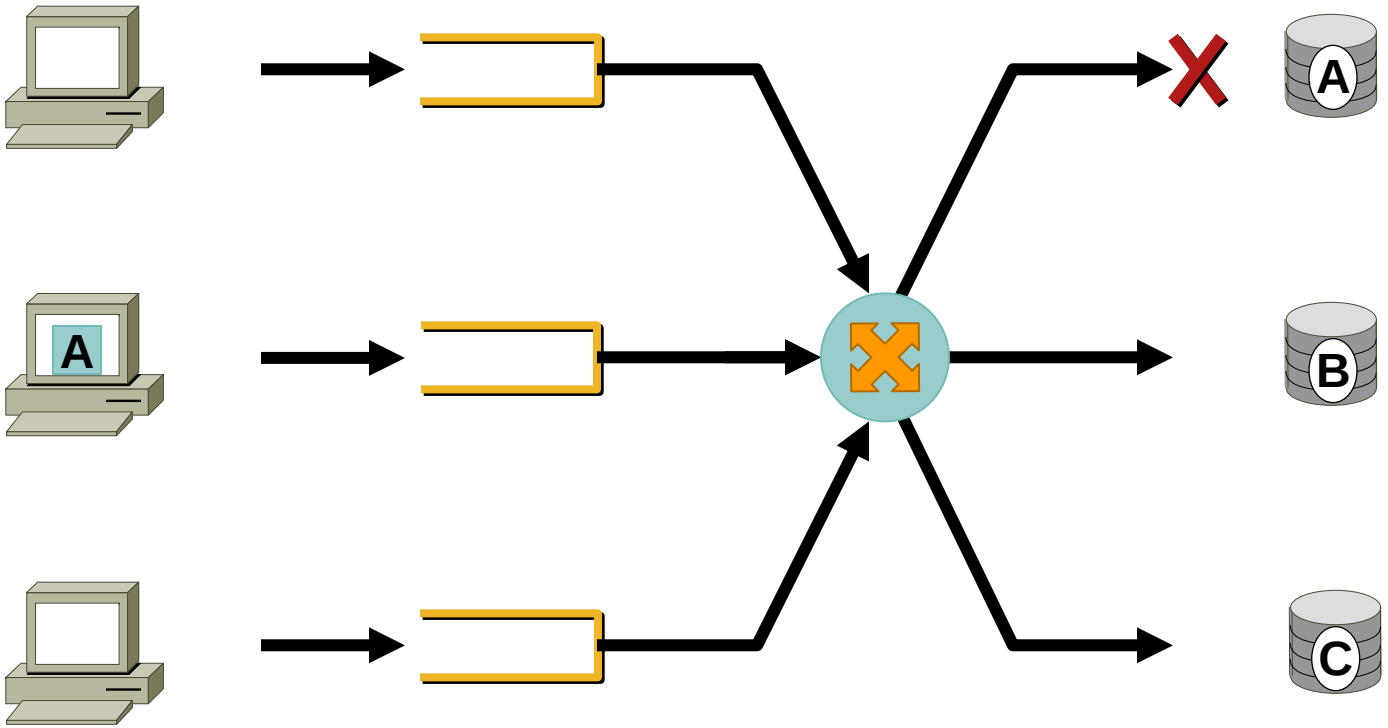

Head of Line Blocking

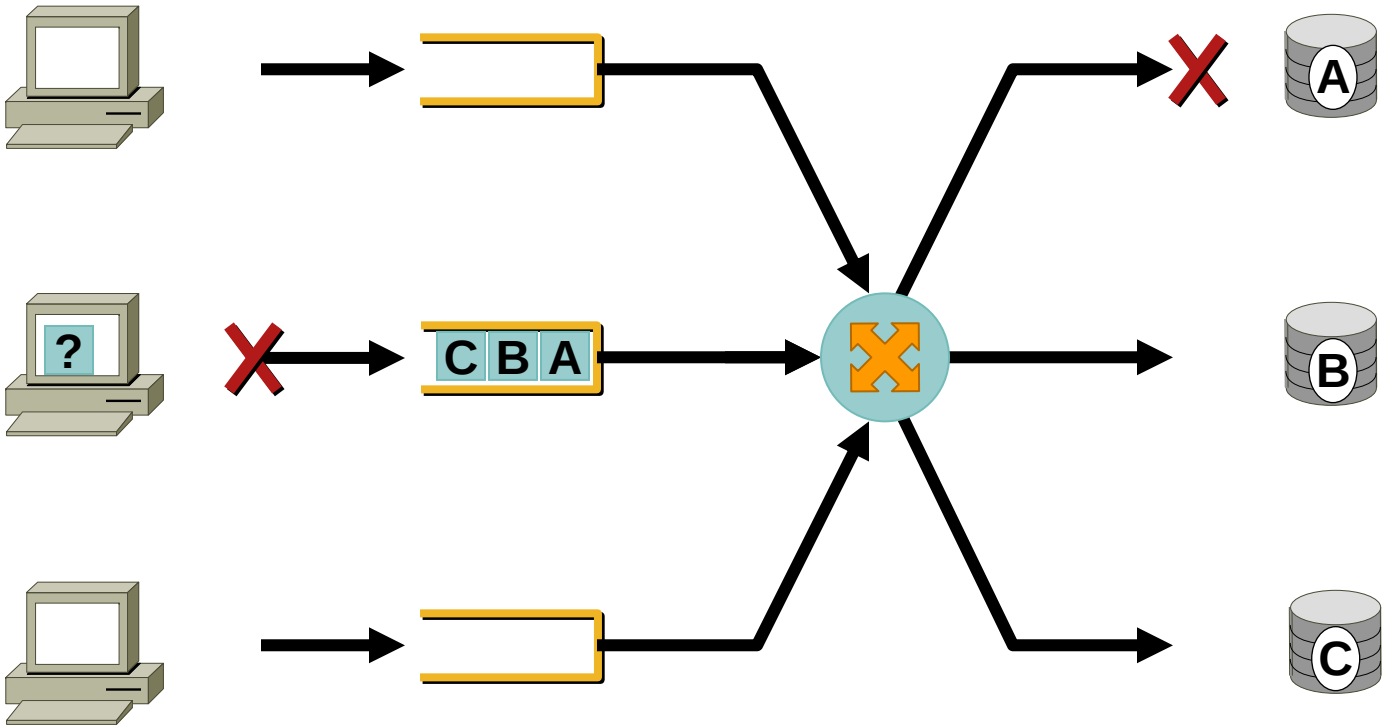

Virtual Output Queues

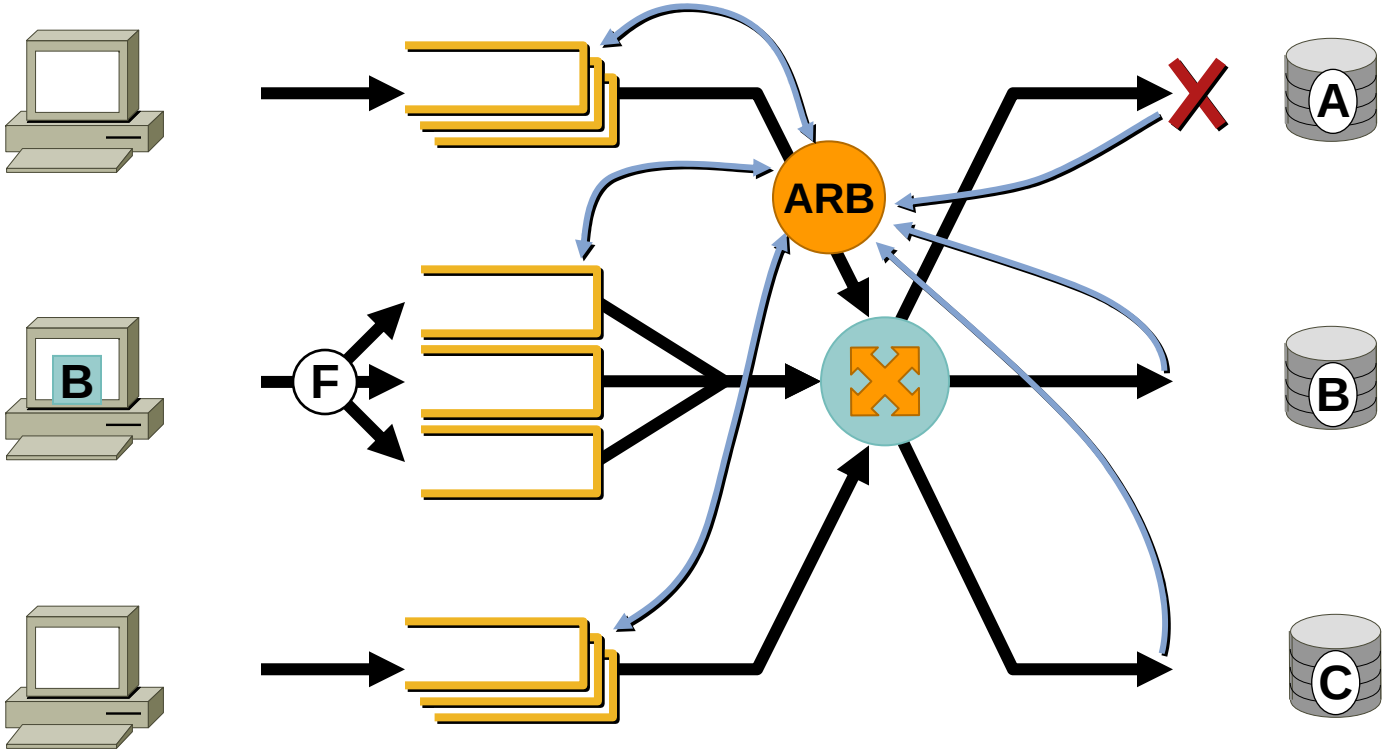

Virtual Output Queues

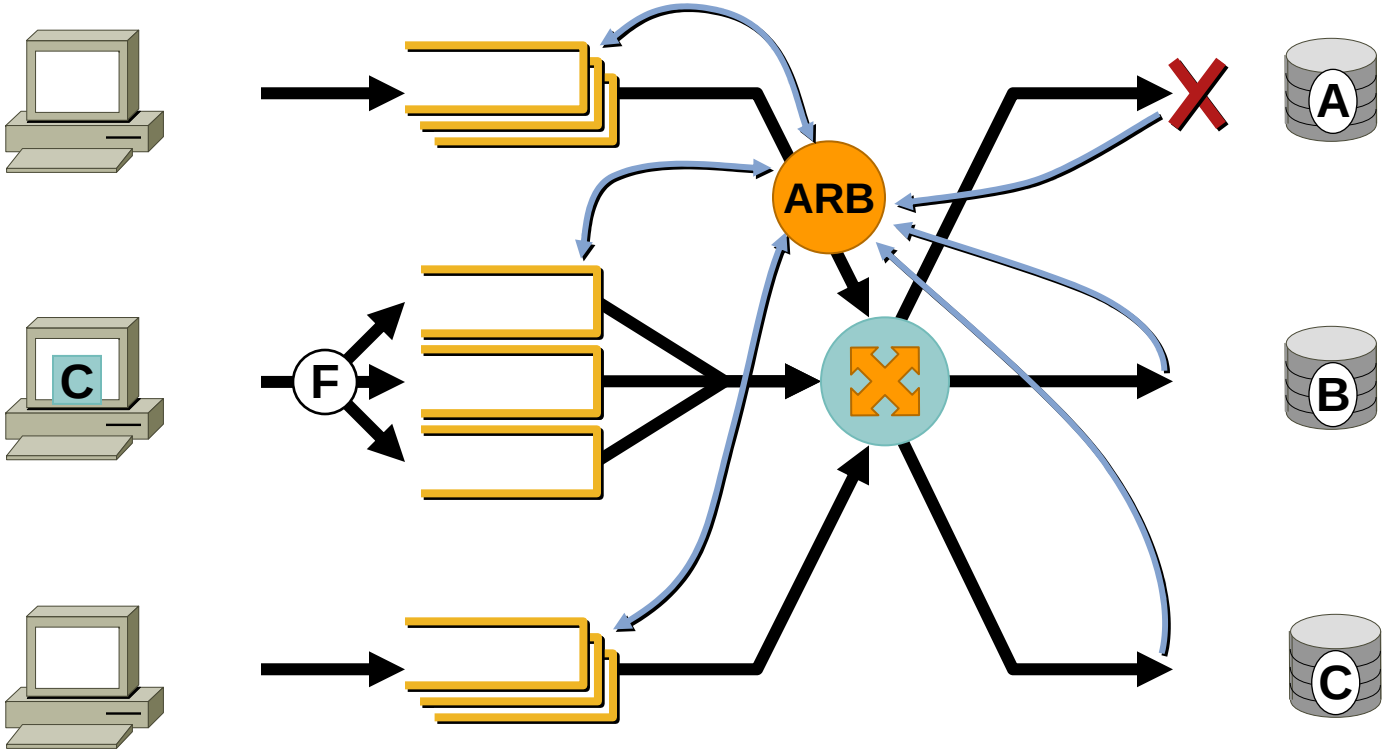

Virtual Output Queues

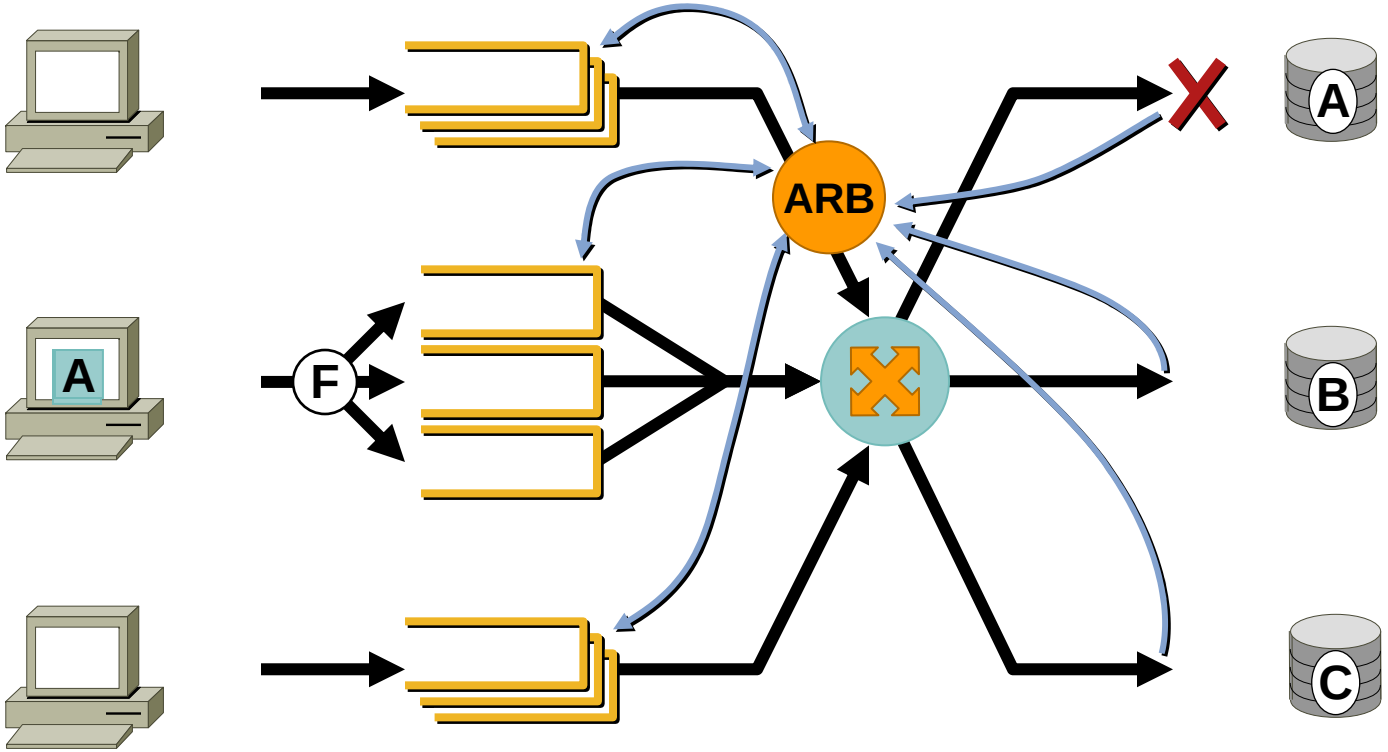

Virtual Output Queues

Avoids Head of line Blocking

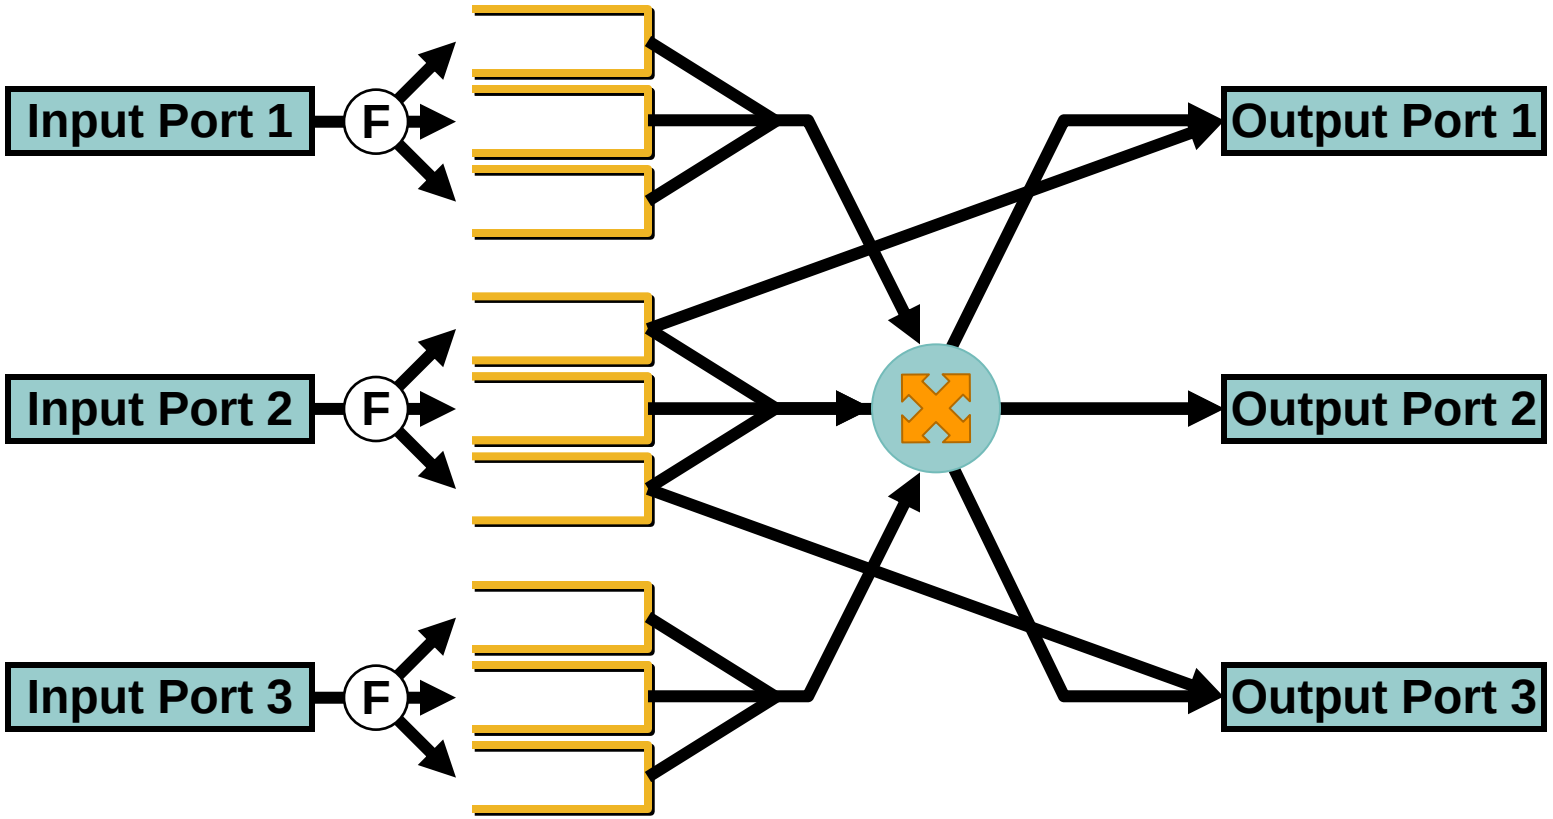

Virtual Output Queues

Larger Effective Buffers

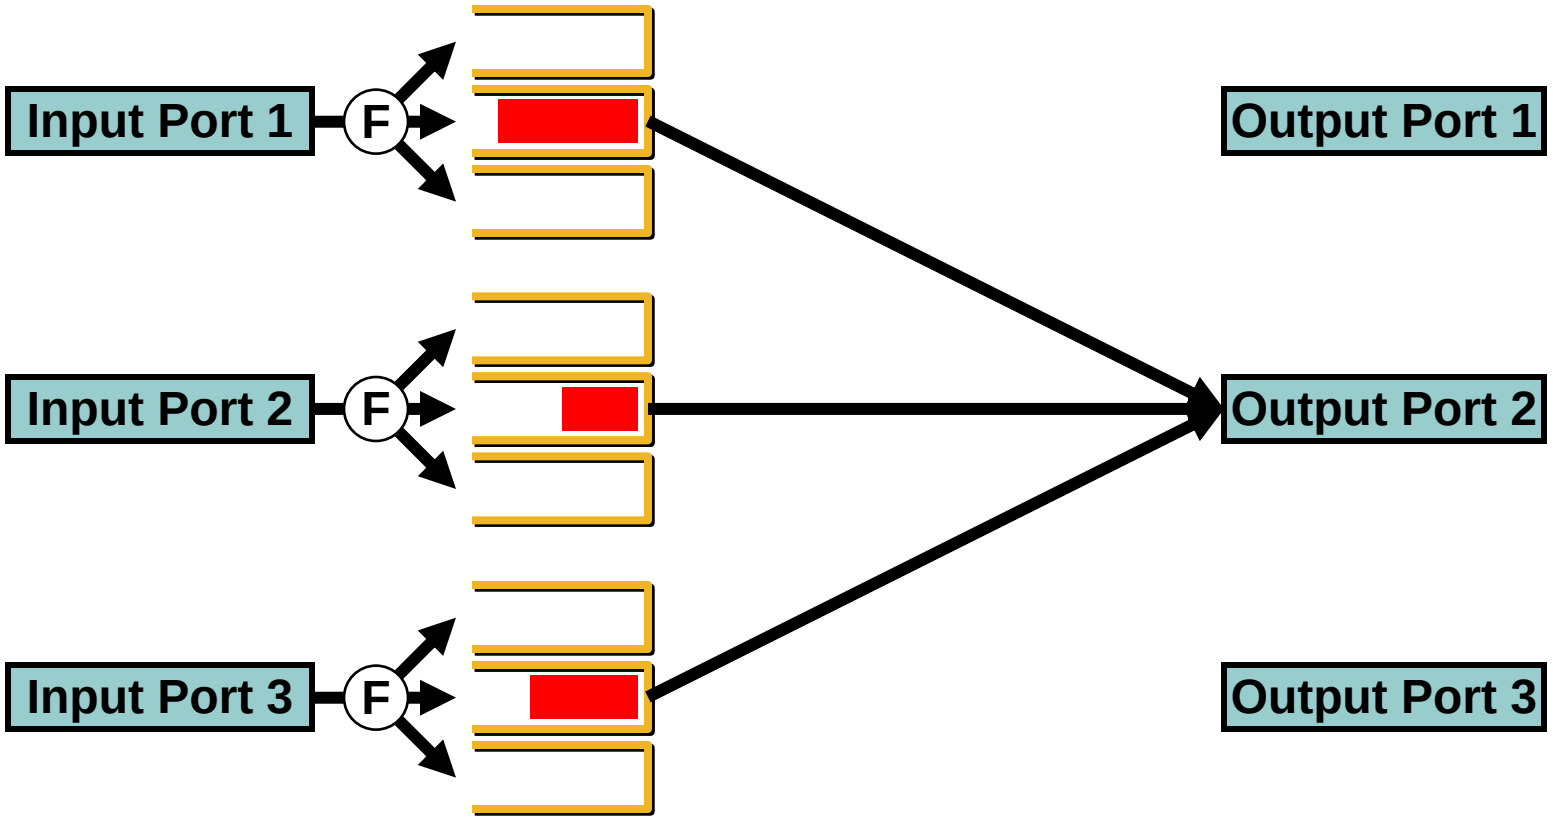

Virtual Output Queues

Better Isolation of Congestion Sources

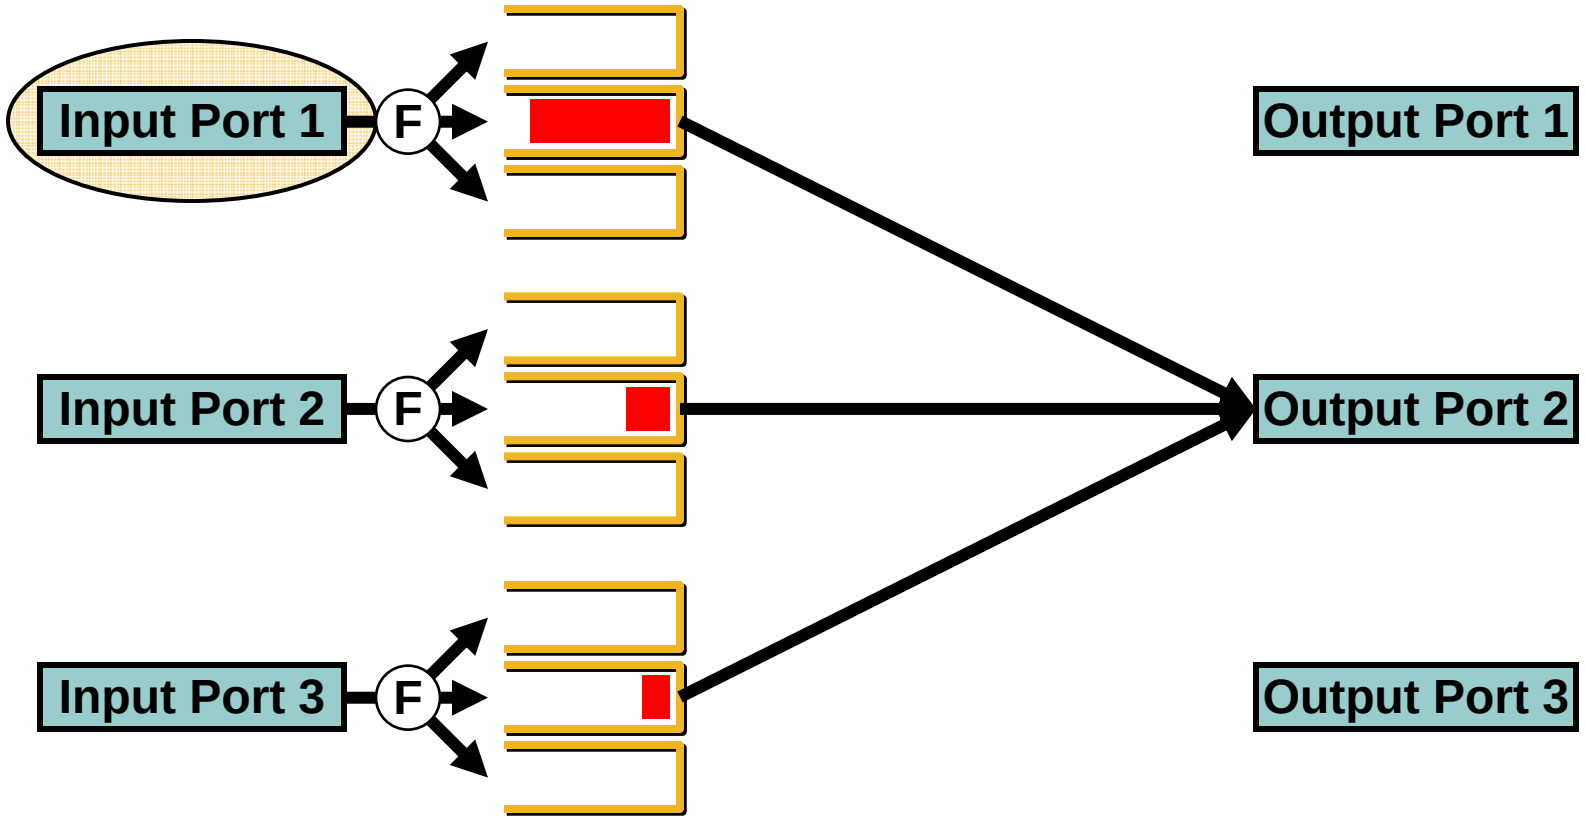

Virtual Output Queues

Better Isolation of Congested Outputs

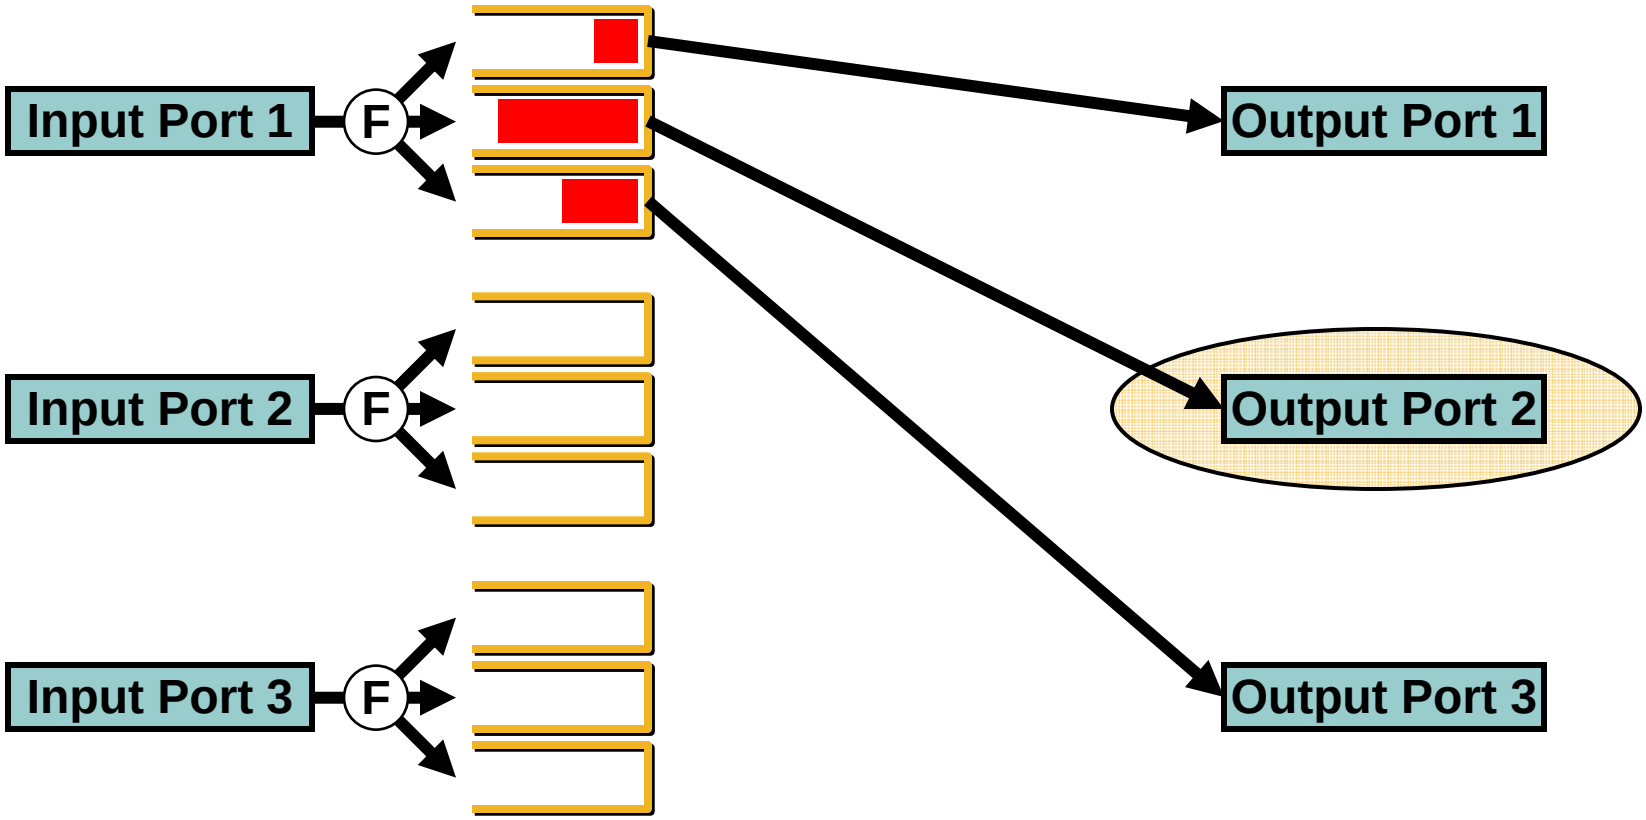

Virtual Output Queues

Better Isolation of Congestion In General

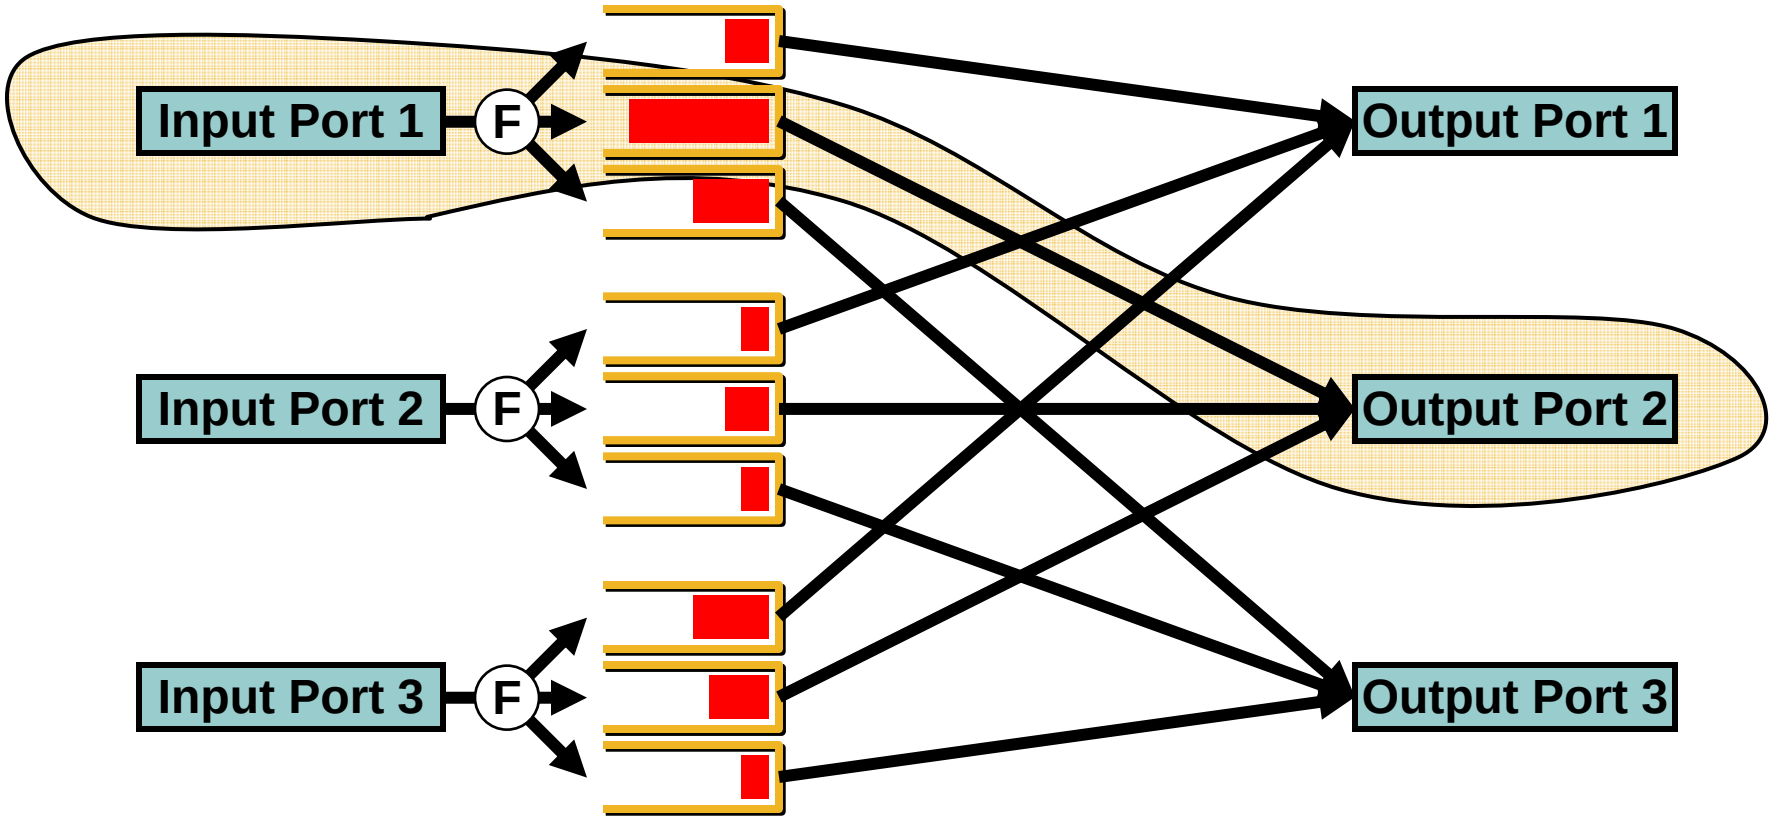

Virtual Output Queues

Buffer Bandwidth Matched to Port Bandwidth

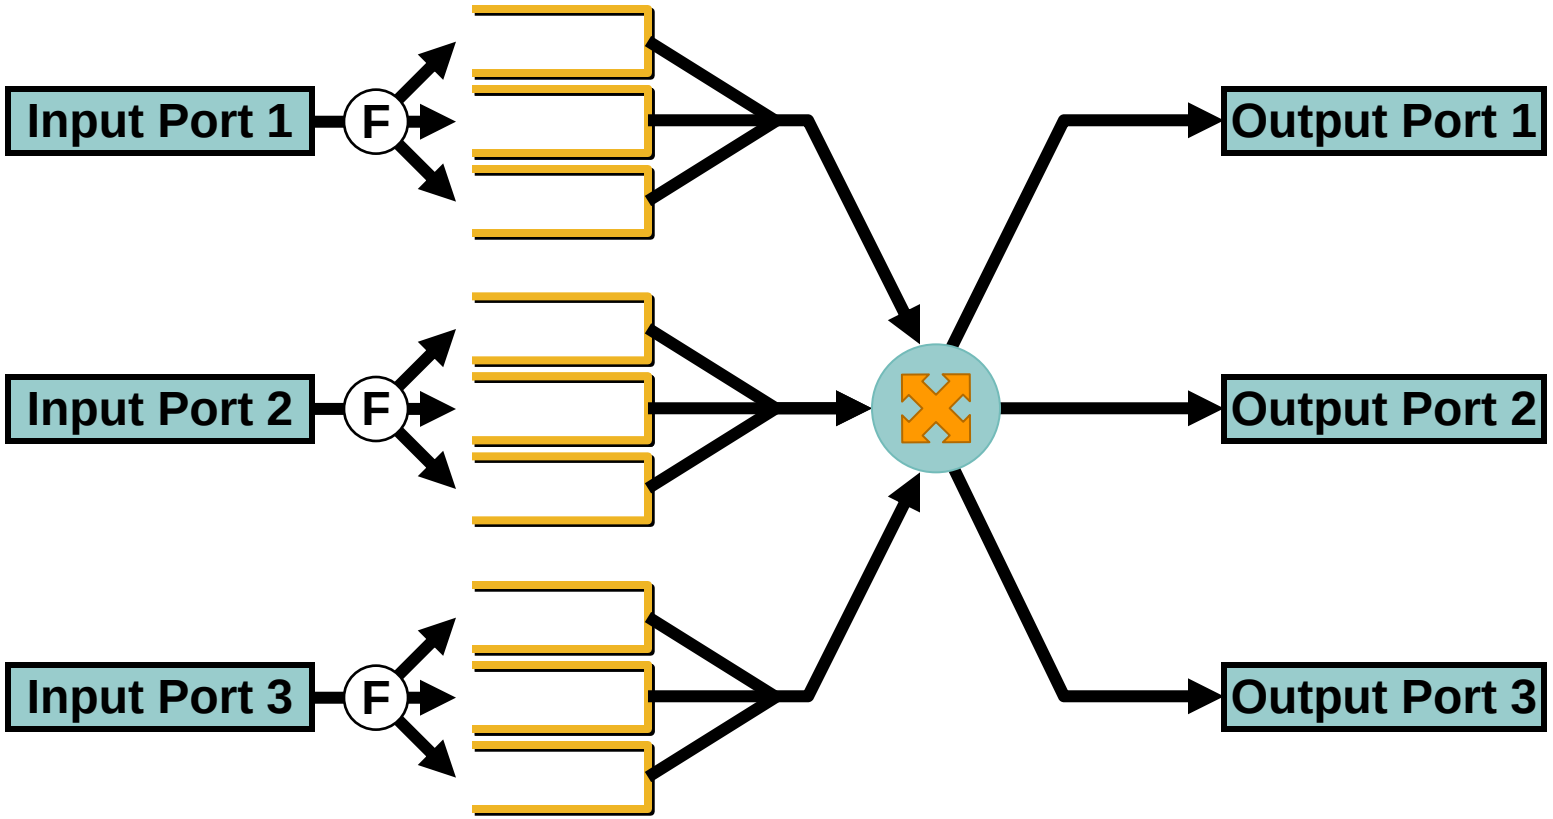

Virtual Output Queuing - Summary

PRO

- Avoids Head-of-line blocking
- More efficient buffer usage
- Better traffic isolation
- Buffer bandwidth matched to port bandwidth
- Necessary for lossless networks.
  (e.g. Fibre Channel, Infiniband, ..)

CON

- Complex to implement
- Locks in scalability

Part III—Real World Implementation

- Cross-bar Performance
  - Classic crossbar vs. Buffered Crossbar
- Virtual Output Queuing
  - Head-of-Line Blocking
- **Completing the Switch**
- Other Features

Completing the Swtich

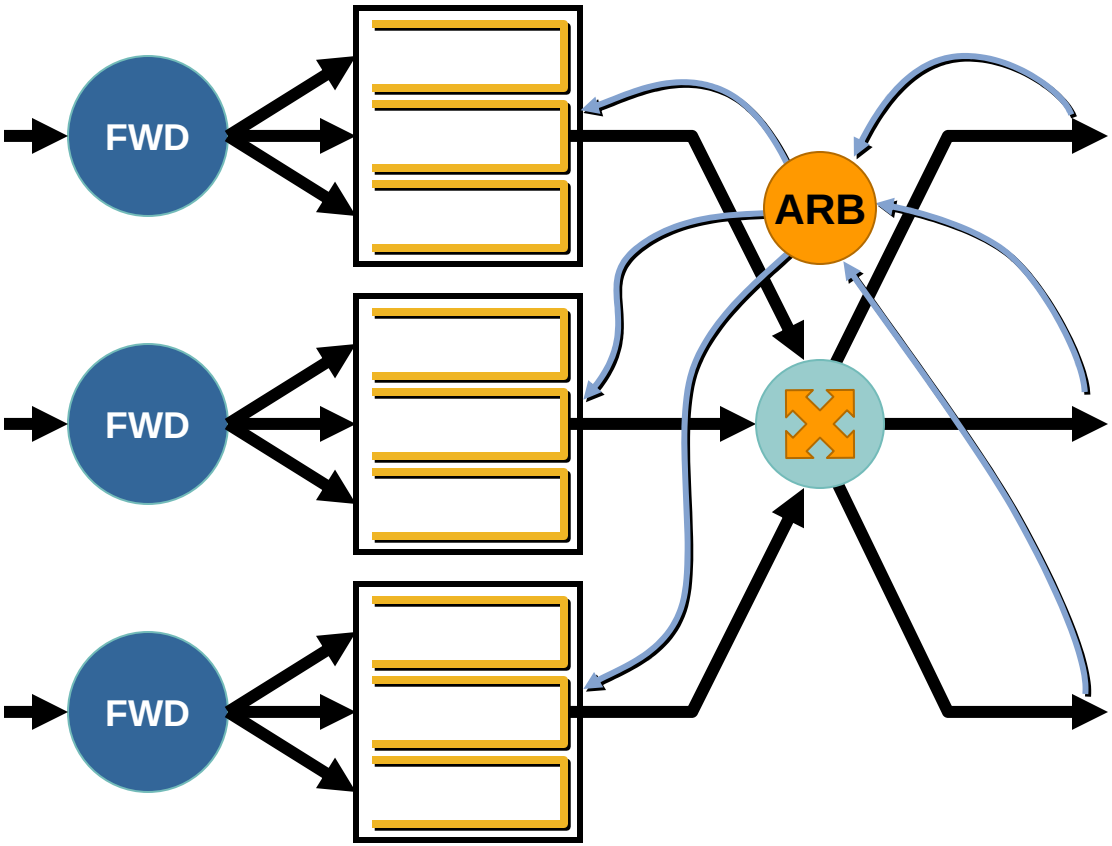

Completing the Swtich

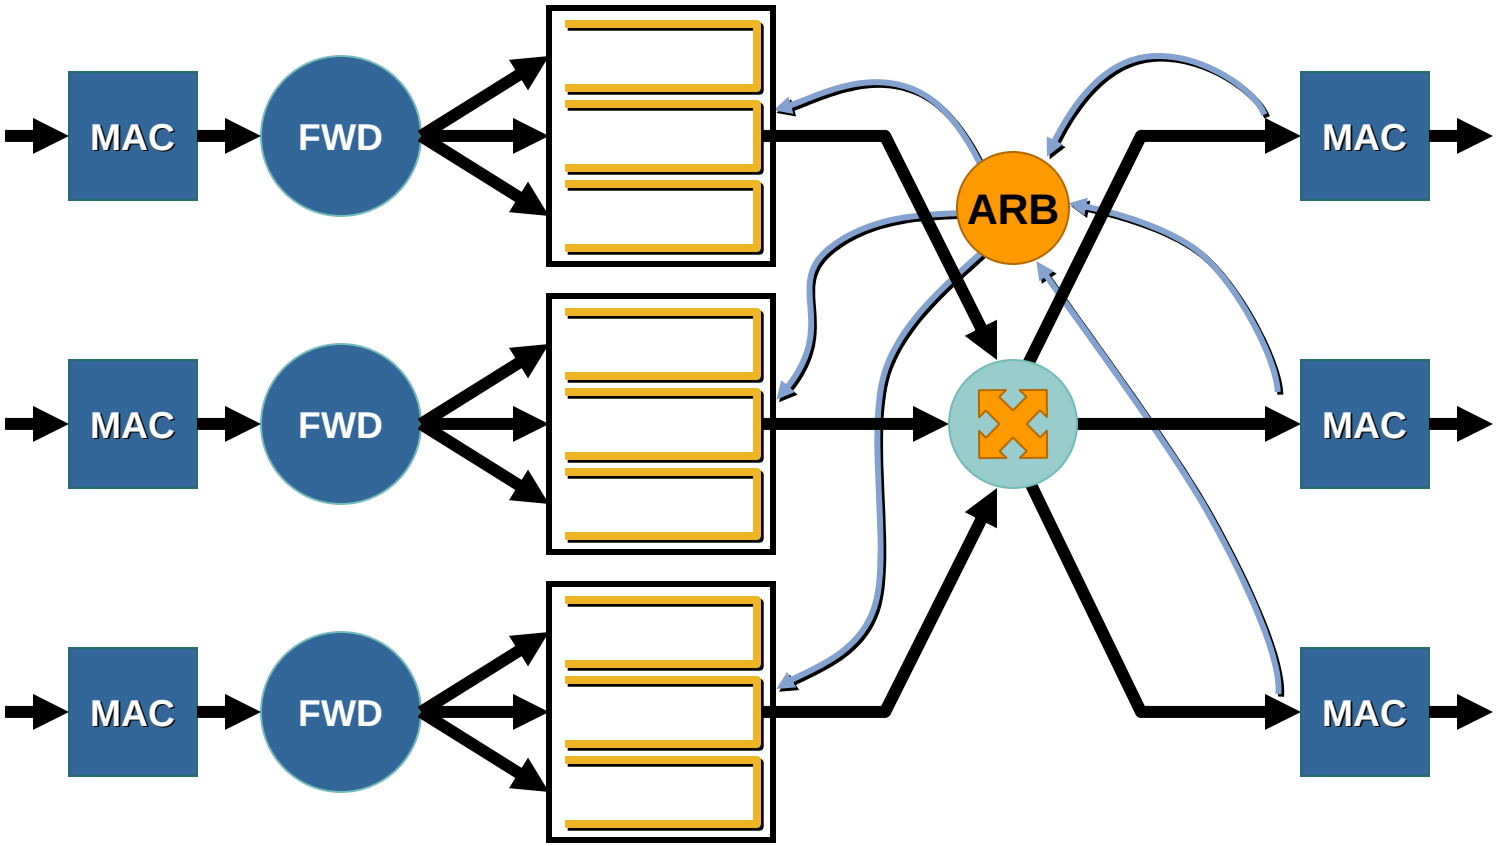

Completing the Swtich

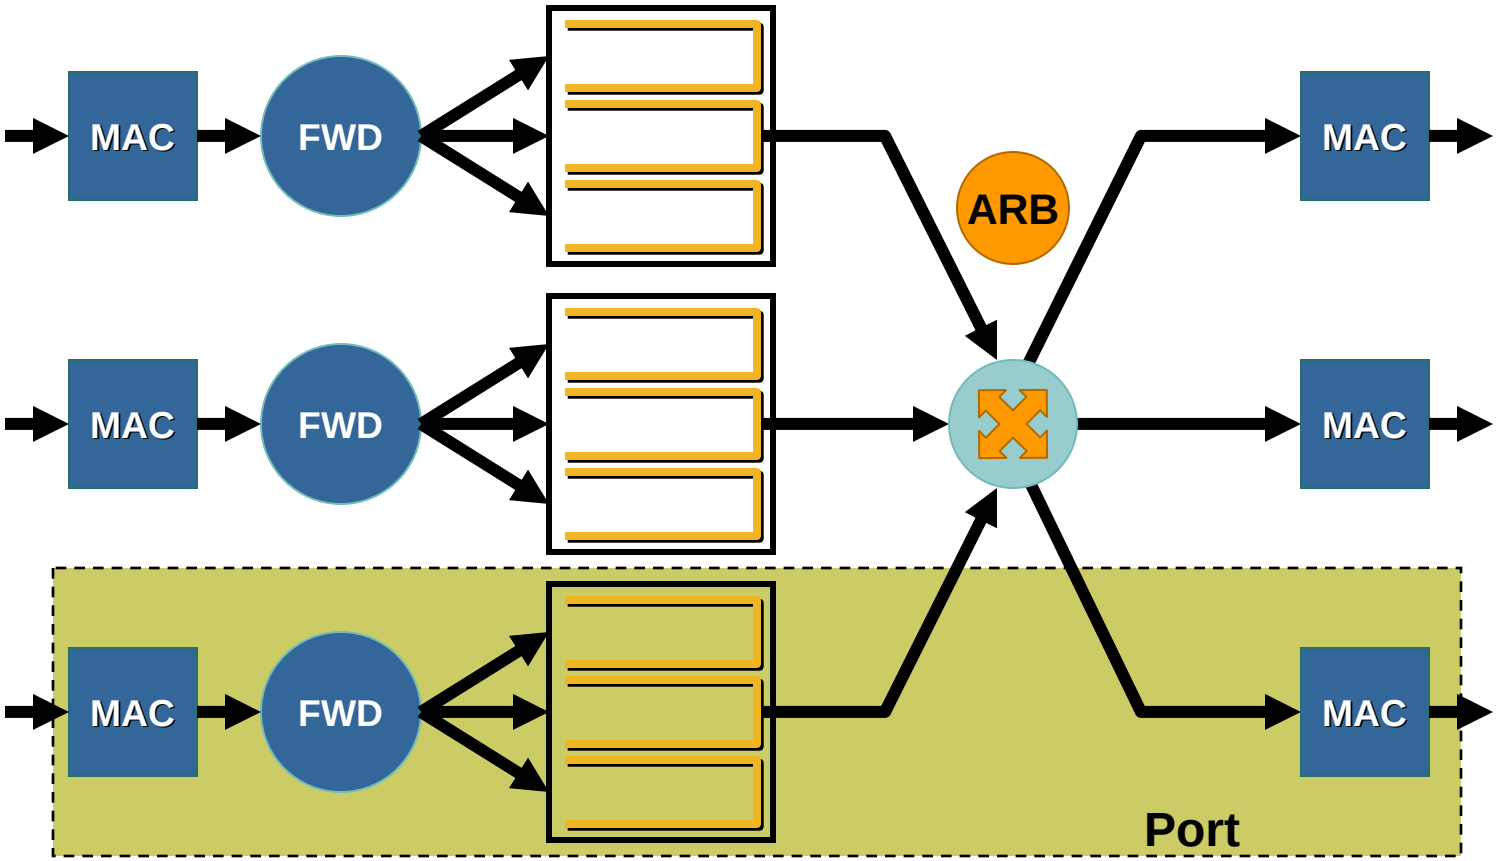

Re-Draw

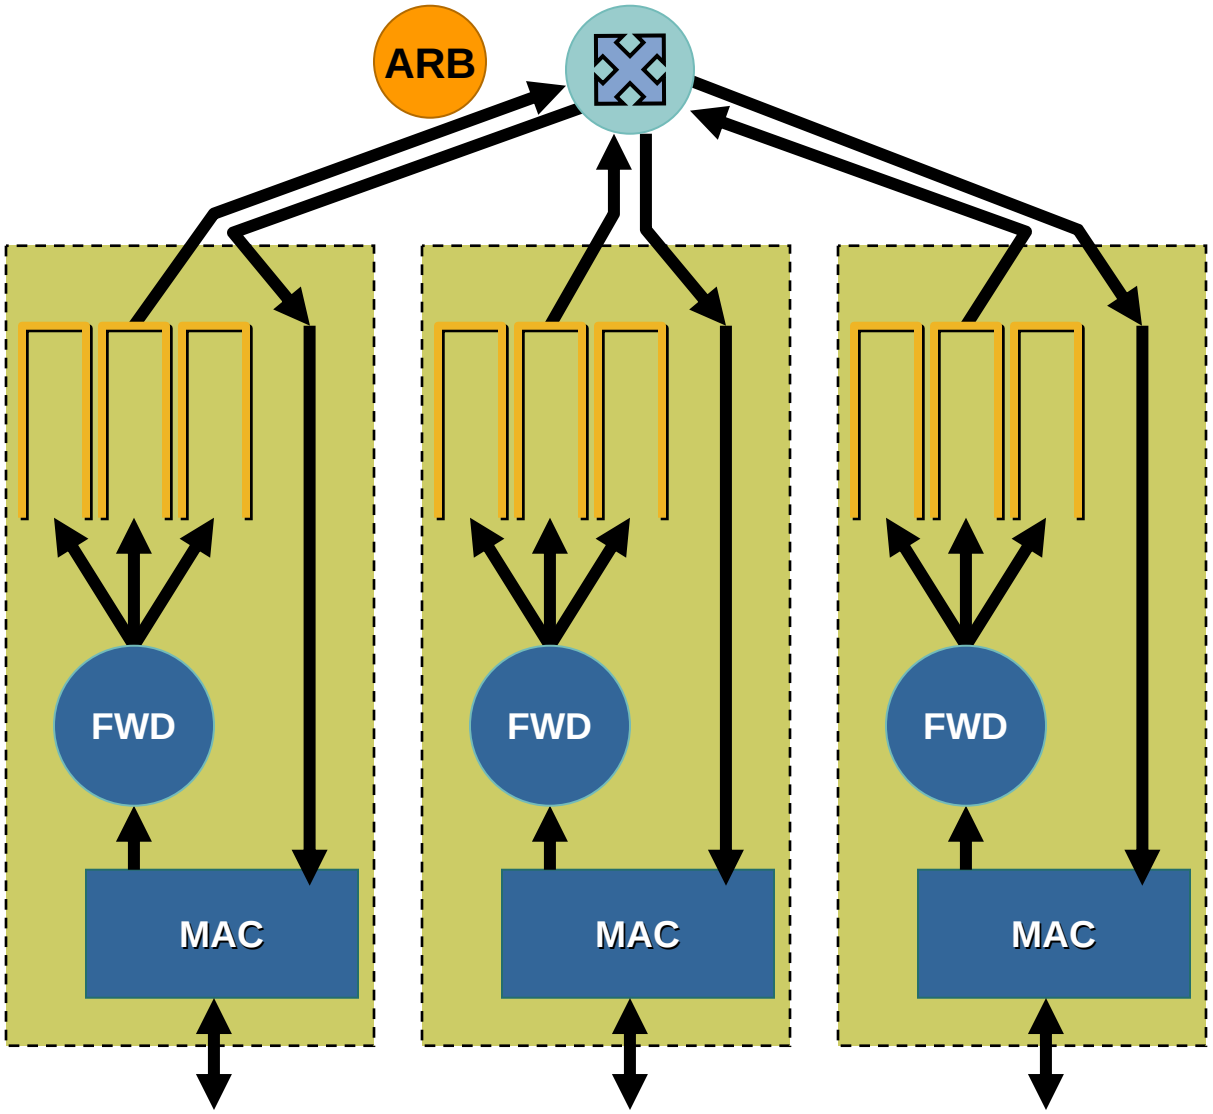

Larger Switch

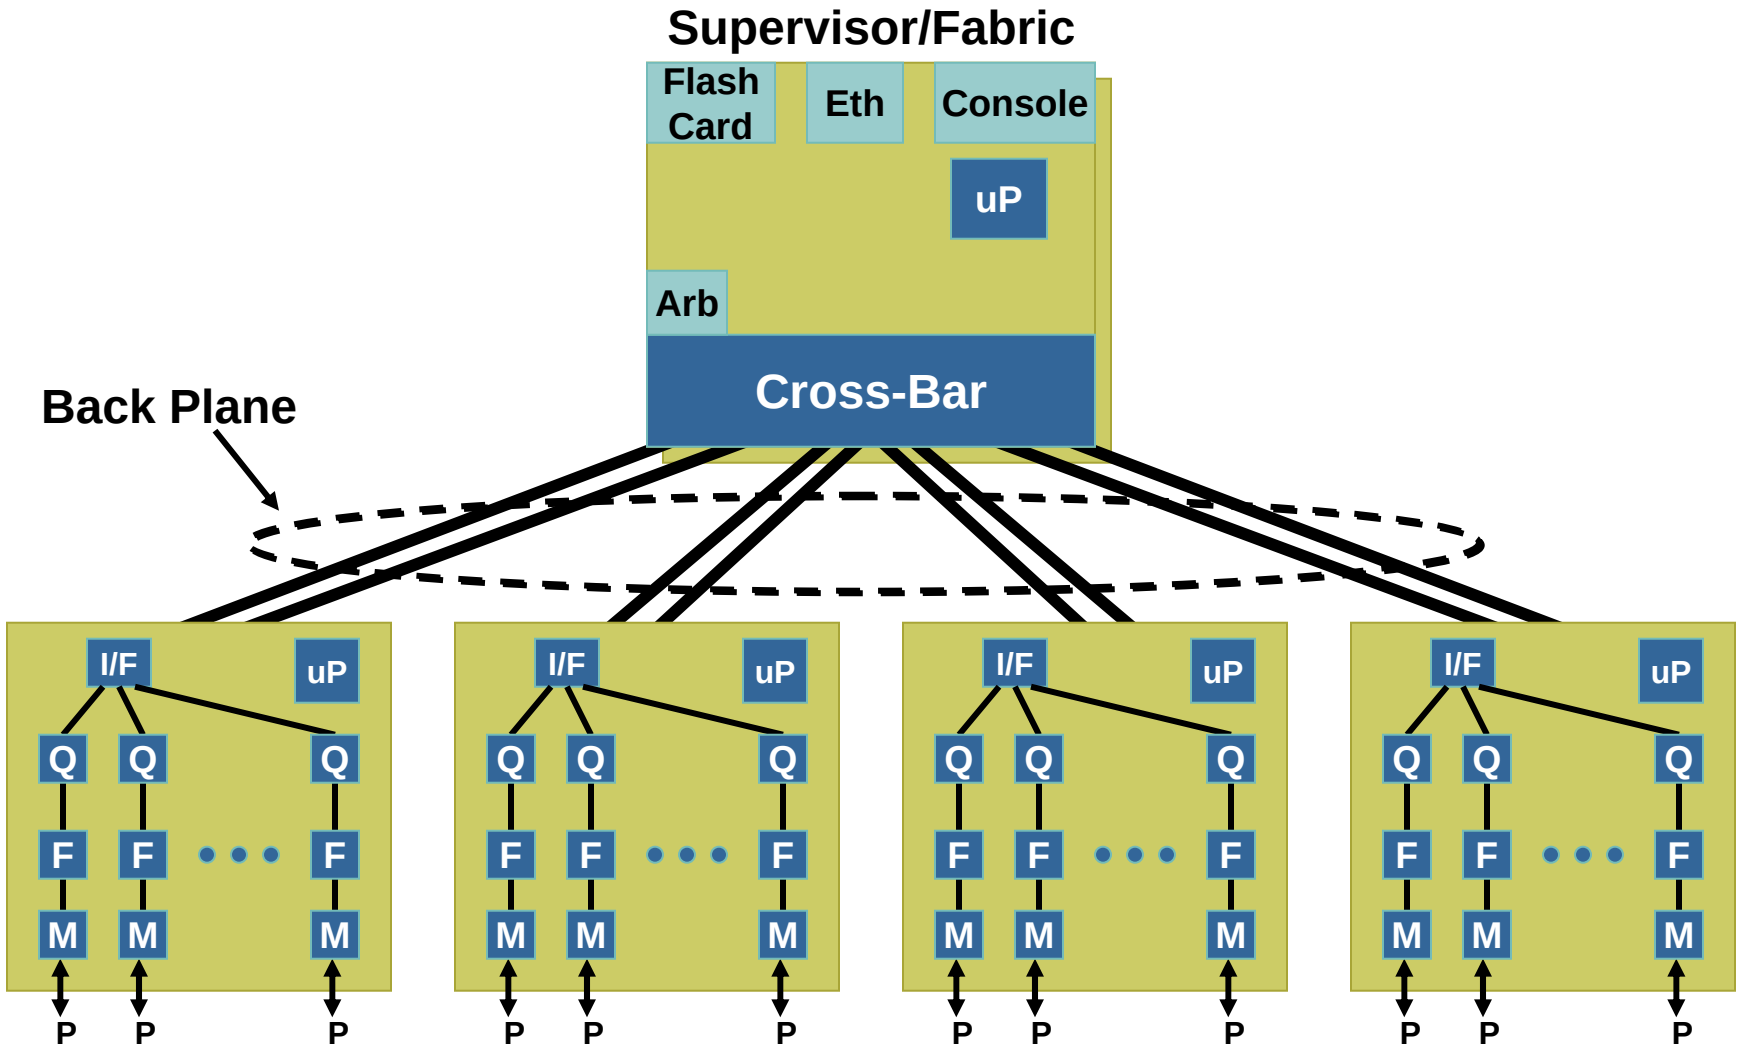

Larger Switch

- Error-Check ingress frame (CRC etc)
- Ingress rate-limiting
- Prepend switch-internal header (Ingress port, Ingress VSAN, QoS bits)
- Timestamp ingress frame arrival time

- Per-VSAN Ingress ACLs (Hard Zoning)
- Per-VSAN Forwarding (Fibre Channel, MPLS)
- Load-balancing decision (determine final egress interface where destination is a PortChannel or equal-cost FSPF route)
- Inter VSAN Routing (Fibre Channel NAT)
- Update src/dst / frame/byte statistics

- Virtual Output Queuing
- Frame Buffering
- Quality of Service (QoS)
- Arbitration and Fairness
- Congestion Control (FCC)

Dual Redundant Centralized Arbiters (capable of scheduling more than 5 billion frames per second)

PHY

Optical to Electrical interface:

- SFP (1 / 2 / 4 Gbps)
- X2 (10Gbps) optics

MAC

Port Speeds:

- 1 / 2 / 4 / 10Gbps

Port Type / Modes:

- F_Port,
- Loop port (FL_Port),
- Private Loop (TL_Port),
- ISL (E_Port),
- EISL (TE_Port),
- SPAN Port (SD_Port)

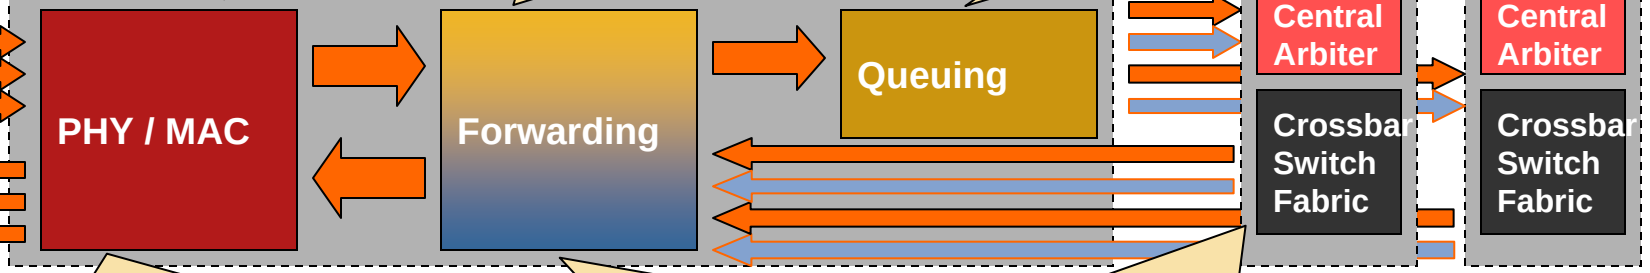

- Transmit, loopback or Expire frame based on ingress timestamp
- Remove switch-internal header from frame
- Prepend EISL header onto frame where egress port is a Trunked E_Port (TE_Port)

- Per-VSAN Egress ACLs (Hard Zoning) including LUN Zoning & read-only Zones
- Re-Validate frame CRC
- Inter VSAN Routing (Fibre Channel NAT)
- Quality of Service (egress QoS)

- Dual Redundant crossbar switch fabrics provide high-speed non-blocking non-oversubscribed capacity from any slot to any slot (25 Gbps raw capacity per channel, 4 channels per linecard)
- Crossbar operates 3X over-speed (75 Gbps) and utilizes superframing to improves crossbar efficiency

Part III—Real World Implementation

- Cross-bar Performance
  - Classic crossbar vs. Buffered Crossbar
- Virtual Output Queuing
  - Head-of-Line Blocking
- Completing the Switch
- **Other Features**

Other Features

- Bridging and Routing
- SVIs
- Multicast
- Broadcast
- ACLs (+VACLs)
- SPAN + RSPAN
- Security (e.g. 802.1x)
- Redundancy
- L4-7 Deep Packet Inspection
- NAT
- MPLS
- *PLS
- GRE
- Virtualization
- Virtual Router Instances
- Services / Service Modules
- Roles / RBAC
- Diagnostics
- (insert here…)

Part IV – The Future

Future Data Center considerations

- Power & Heat becoming the primary consideration

  Density of servers rapidly shrinking;
  Server density + Server capacity typically growing >3x faster than associated Cooling/AC infrastructure

- Moore's Law

  Today's Commodity CPUs were yesterday's Supercomputers

- Virtualization

  Server, Network, Storage, …

- Future Data Center Network

  At most, dual server connections

  Higher Speeds, Lower Latency

  100 Mbps ▸ GbE ▸ 10GbE ▸ 100+GbE?

  Infrastructure Simplification

Future Data Center considerations

- **Power & Heat becoming the primary consideration**

  Density of servers rapidly shrinking;
  Server density + Server capacity typically growing >3x faster than associated Cooling/AC infrastructure

- Moore's Law

  Today's Commodity CPUs were yesterday's Supercomputers

- Virtualization

  Server, Network, Storage, …

- Future Data Center Network

  At most, dual server connections

  Higher Speeds, Lower Latency

  100 Mbps ▸ GbE ▸ 10GbE ▸ 100+GbE?

  Infrastructure Simplification

Data Center Physical Infrastructure

Desire is to Last 10 to 15 Years

Useful life of a Data Center facility is expected to exceed the life of the gear it contains due to the tremendous cost of relocating or renovating data center facilities

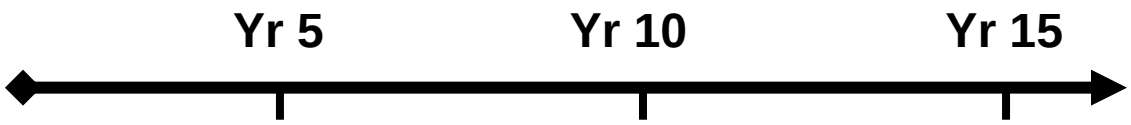

5 to 7 Server Upgrade Cycles

Mainframes → 1/2RU Servers → Multi-Core CPU's → Blade Server

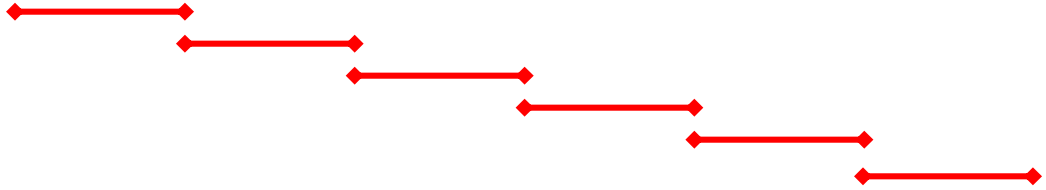

2 to 3 Network Infrastructure Upgrades

10/100→10/100/1000 with L4-7 Services→ 10GbE Attached hosts

Obtaining 10-15 Years of Useful Life Requires Planning for Growth of Power and Cooling

Planning and Coordination Between IT Staff and Facilities Managers is a Must

2kW/Rack → 6kW/Rack → 15+kW/Rack!

5 to 7 Server Upgrade Cycles

Mainframes → 1/2RU Servers → Multi-Core CPUs → Blade Server

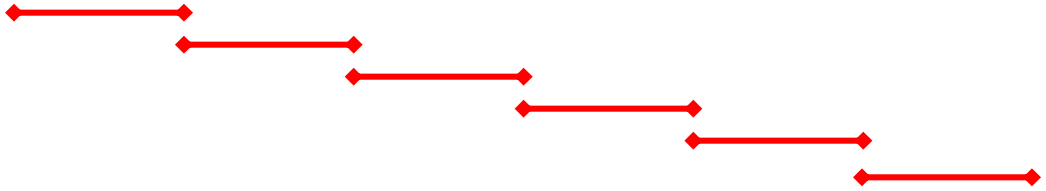

2 to 3 Network Infrastructure Upgrades

10/100 → 10/100/1000 with L4-7 Services → 10GbE Attached hosts

Converged LAN/SAN/IPC

Future Data Center considerations

- Power & Heat becoming the primary consideration

  Density of servers rapidly shrinking;
  Server density + Server capacity typically growing >3x faster than associated Cooling/AC infrastructure

- **Moore's Law**

  Today's Commodity CPUs were yesterday's Supercomputers

- Virtualization

  Server, Network, Storage, …

- Future Data Center Network

  At most, dual server connections

  Higher Speeds, Lower Latency

  100 Mbps ▸ GbE ▸ 10GbE ▸ 100+GbE?

  Infrastructure Simplification

Moore's Law

An empirical observation

"The transistor density of Integrated Circuits, with respect to minimum component cost, will double every 24 months."

— Gordon E Moore, co-founder, Intel, 1965

- 1978: 8088: 20K transistors
- 1985: Intel® 386™: 275K transistors
- 1989: Intel® 486™: 1.2M transistors
- 1993: Intel® Pentium™: 3.1M transistors
- 1999: Intel® Pentium 3™: 9.5M transistors
- 2000: Intel® Pentium 4™: 42M transistors
- 2006: Intel® Core™2 Duo: **291M transistors**

Compute Capacity is Increasing
Network Challenges & Implications

- Growth is not slowing down
- More effective way to scale layer 2 networks is needed
- Ready access to compute AND storage resources
- Powering 10GbE over Copper PHYs on high density linecards (40+ ports)
- Impact of an outage/downtime will be magnified- maximize *operational* uptime and cluster utilization
- Least disruptive upgrade path – hardware, software & services

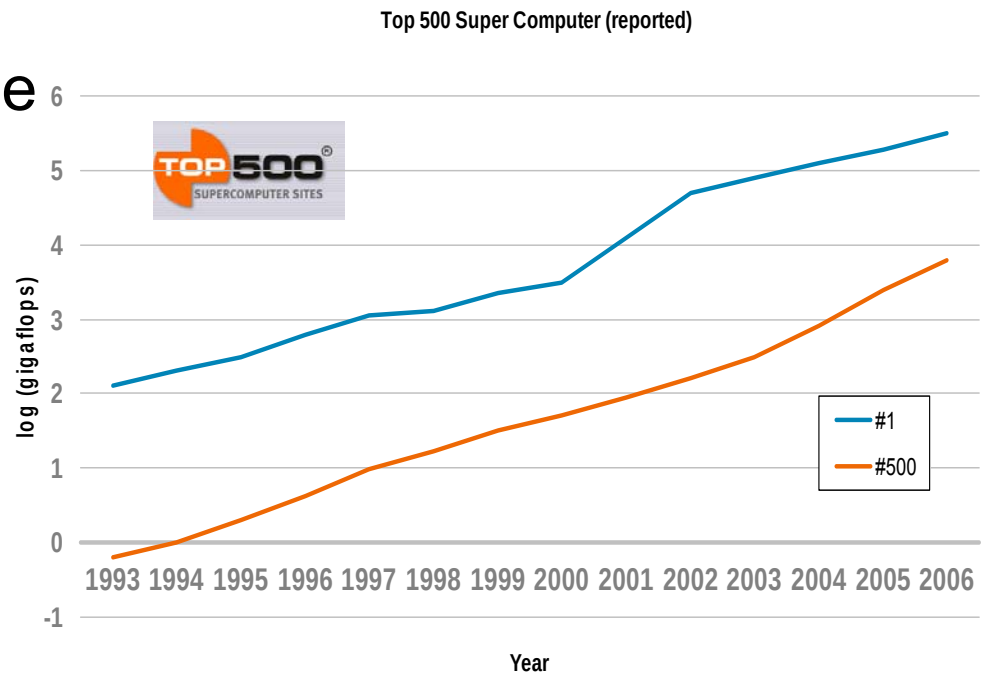

Future Data Center considerations

- Power & Heat becoming the primary consideration

  Size of servers rapidly shrinking;
  Server density + Server capacity typically growing >3x faster than associated Cooling/AC infrastructure

- Moore's Law

  Today's Commodity CPUs were yesterday's Supercomputers

- **Virtualization**

  Server, Network, Storage, …

- Future Data Center Network

  At most, dual server connections

  Higher Speeds, Lower Latency

  100 Mbps ▸ GbE ▸ 10GbE ▸ 100+GbE?

  Infrastructure Simplification

VIRTUALIZATION

The capability to decouple logical functions from physical devices.

- **Increase resource usability –**
  Lower CAPEX
- **Loosely couple re-usable functions –**
  Flexibility
- **Dynamic allocation of virtual instances –**
  Automation
- **Centralized policy Management –**
  Lower TCO
- **Distributed Capabilities –**
  Broad use of functions

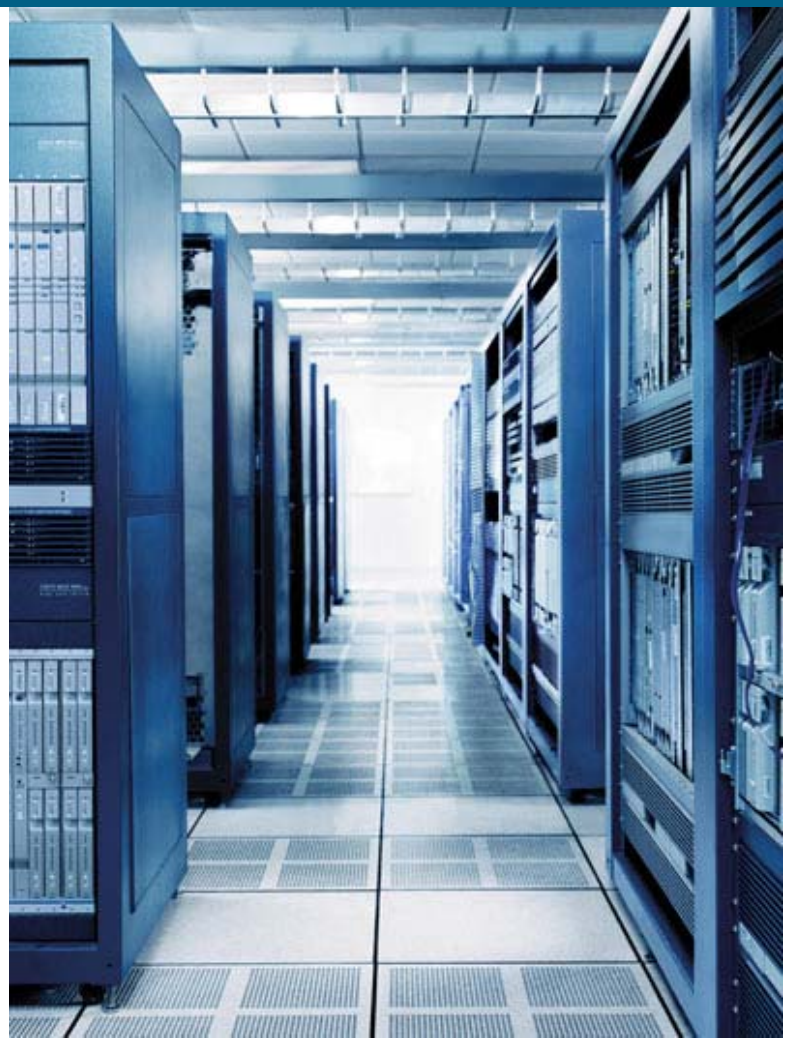

Data Center Virtualization
Examples

Type	Example	Benefit
Pooling	802.3ad, Port Channel TCP Connection Pooling Symmetric Multi-processing Oracle Cache Fusion	Improved availability Cost effective scaling Simplify topologies and configuration
Partitioning	VLAN, VSAN Virtual Firewall, ACE VMWare	Increased utilization Reduce number of devices Segmentation of policy/ traffic
Abstraction	HSRP/VRRP, GLBP, WWNS VIP/NAT vFrame	Simplify provisioning *Provision ACROSS technology silos*

Evolution in Data Center Computing

Compute
LAN
SAN

Mainframe and SMP	x86 uprocessor	Clusters/ multi-core CPU's	Virtual Machines	Utility Computing
• Compute, Storage and Interconnect Resources *Physically* Integrated • Few Applications • Proprietary Architecture • Minimal Connectivity Requirements	• Reduction Apps Development Cost • 1st Gen CPU's not very powerful • Client/Server Model	• PCI, PCIe, PCIx • Emergence of Blade Servers • Compute nodes increasingly more powerful	• Location of Servers Dynamic • Multiple Apps per CPU	• Compute, Storage and Interconnect Resources *Logically* Integrated • Geographic Separation of Resources • Large # of Applications • On Demand Service Provisioning
	• Clustering Technologies • TCP/SSL Offload • VLAN's and L3 • VTP and VLAN pruning • ACL's and Security	• Myrinet and IB as extension of PCI • QoS • 10/100 -> 10/100/1000 -> 10GbE	• Large, flat L2 domains • L4-7 service chaining (ACE) • Services Centric Provisioning Model (vFrame)	
	• DAS migration to NAS, SAN, FC-SAN, VSAN • iSCSI	• 4G and 10G FC • Parallel File Systems with SAN Backend • iSLB	• VSANs with Hard-zoning • Dynamic SAN environment	

Future Data Center considerations

- Power & Heat becoming the primary consideration

  Density of servers rapidly shrinking;
  Server density + Server capacity typically growing >3x faster than associated Cooling/AC infrastructure

- Moore's Law

  Today's Commodity CPUs were yesterday's Supercomputers

- Virtualization

  Server, Network, Storage, …

- **Future Data Center network**

  At most, dual server connections

  Higher Speeds, Lower Latency

  100 Mbps ▸ GbE ▸ 10GbE ▸ 100+GbE?

  Infrastructure Simplification

At most, dual server connections

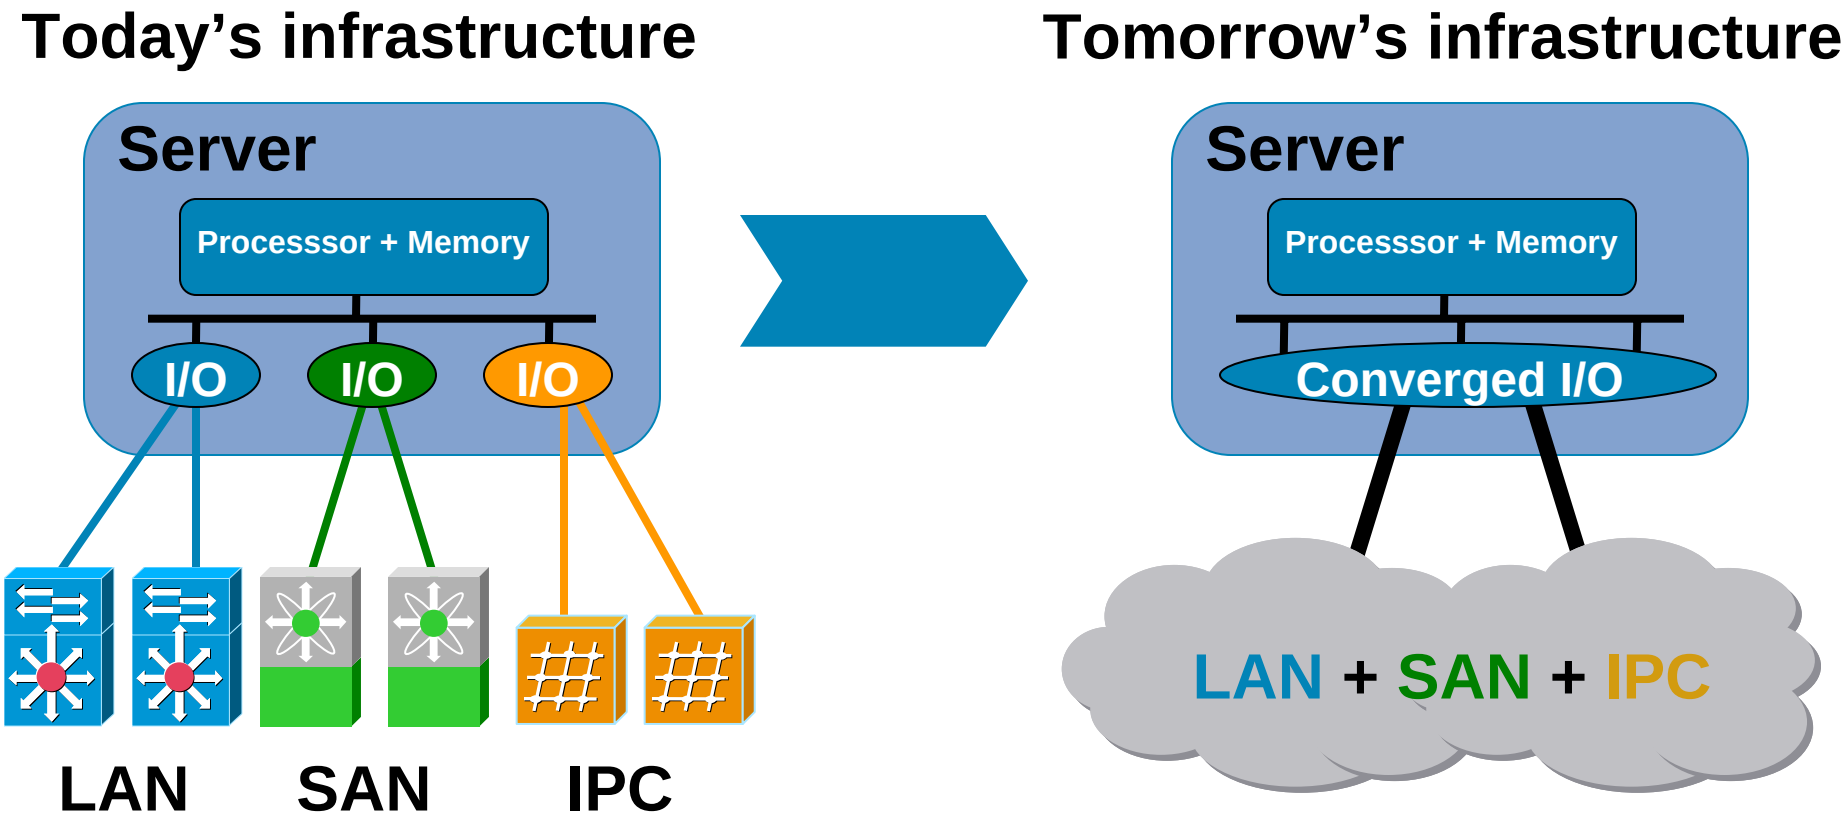

- I/O Consolidation simplifies platform architectures, reducing overall cost

Growing demands of Ethernet

- **1995**: Catalyst 5000 introduced:
  - **16K MAC** addresses
  - 192KB buffering / 8 ports (24KB/port) (10Mbps, 10/100Mbps)
- **1999**: Catalyst 6000 introduced:
  - **128K MAC** addresses (32K-80K usable)
  - **512KB/port** buffering (GbE port)
- **2003**: Catalyst 6500 Series 4 x 10GbE modules
  - **18MB/port** buffering (10GbE ports)
- **2006**: Catalyst 6500 Series 8 x 10GbE module
  - **200MB/port** buffering (10GbE ports)
  - Metro Ethernet environments demanding **>128K MAC addresses**

Growing demands of Ethernet

- CAM tables (L2), ACLs, QoS, CEF FIB tables (L3), Statistics all typically built using Tenary CAMs
- Tenary CAM ≈ Static RAM (SRAM)
  - –SRAM requires 6 transistors per bit
  - –'always ON' (high power)
- Transistor count for 128K MAC addresses, CEF FIB, QoS, Statistics, ...

  now typical to have over 200M transistors just in tables
- comparison:

  Intel® Core™2 Duo "Extreme Edition":
  291M transistors

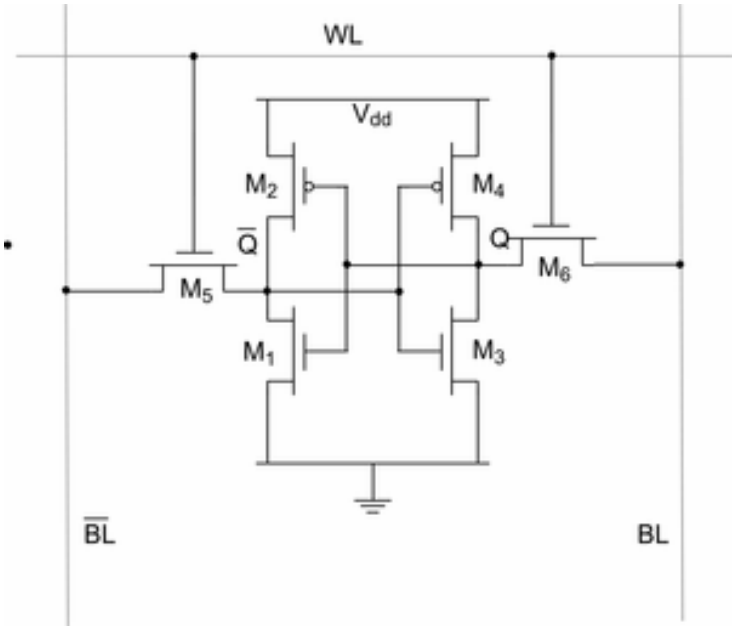

A six-transistor CMOS SRAM cell.

source: Wikipedia
http://en.wikipedia.org/wiki/Static_random_access_memory

Moore's Law

An empirical observation

"The transistor density of Integrated Circuits, with respect to minimum component cost, will double every 24 months."

— Gordon E Moore, co-founder, Intel, 1965

- 1978: 8088: 20K transistors
- 1985: Intel® 386™: 275K transistors
- 1989: Intel® 486™: 1.2M transistors
- 1993: Intel® Pentium™: 3.1M transistors
- 1999: Intel® Pentium 3™: 9.5M transistors
- 2000: Intel® Pentium 4™: 42M transistors
- 2006: Intel® Core™2 Duo: **291M transistors**

Programmability

- **Cisco CRS-1 Carrier Routing System**
- Heart of each linecard is the “Metro” ASIC
- Programmable 40Gbps Network Processor
  - 192 Cores per chip
  - Fully programmable
  - PowerPC-like instruction set

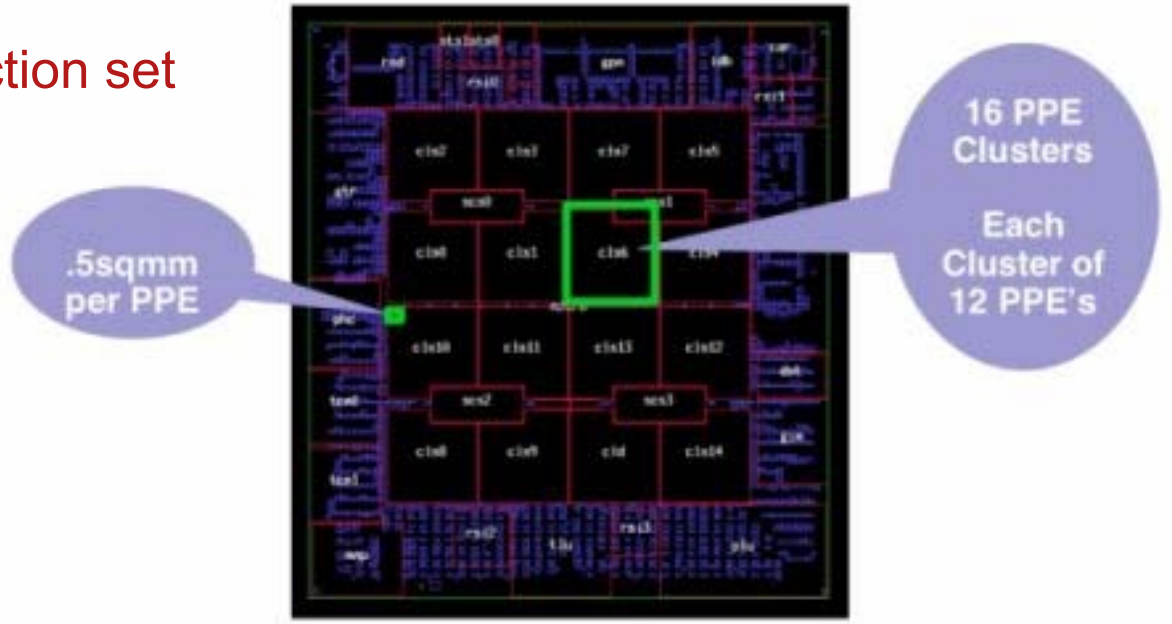

Cisco Data Center Network Design Goals

- Present a Unified Fabric Comprising Multiple Technologies. Service vs Technology Virtualization
- Increase Service/ Applications Deployment Velocity while Reducing Marginal Costs
- Enable Layer 2 Networks to continue to scale beyond spanning tree limitations
- Non-Intrusively Improve Applications Performance and Availability
- Reduce the Number of Audit Points
- Enable Per Business Unit and Per Application SLA's
- Support 24x7x365 operations

Q and A
Questions?