



Server Use Cases for 40GE and 100GE

July 11, 2007

Robert Winter, Dell
robert_winter@dell.com

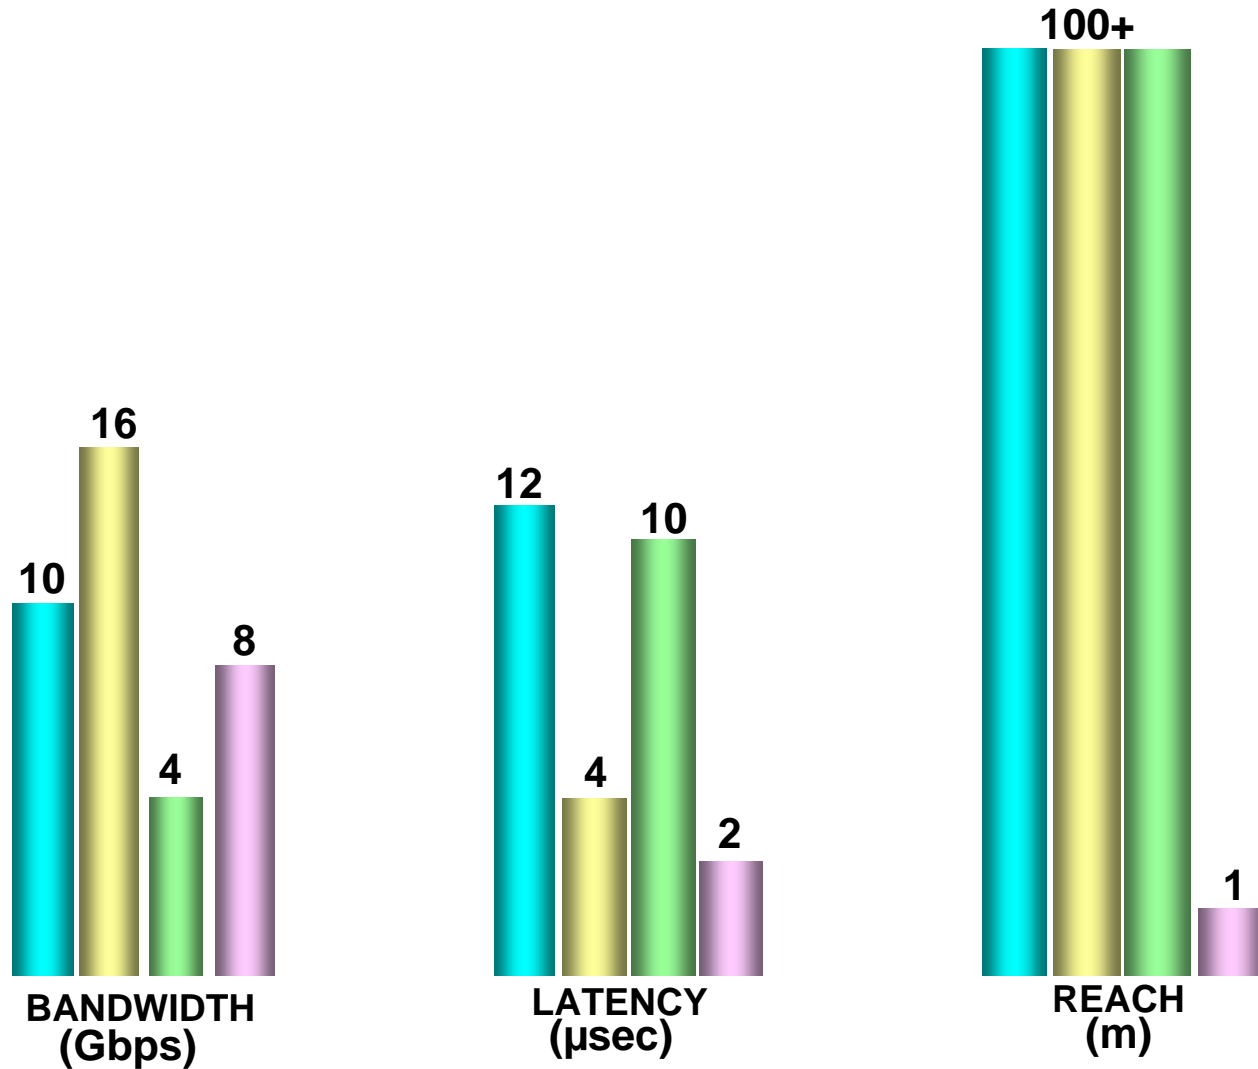


Fabric Comparison

10 Gb/s Ethernet is “good enough”
(for all applications)

- EN, 10Gbps
- IB, 4xDDR
- FC, 4
- PCIe gen1 x4

2007
(actual data rates)



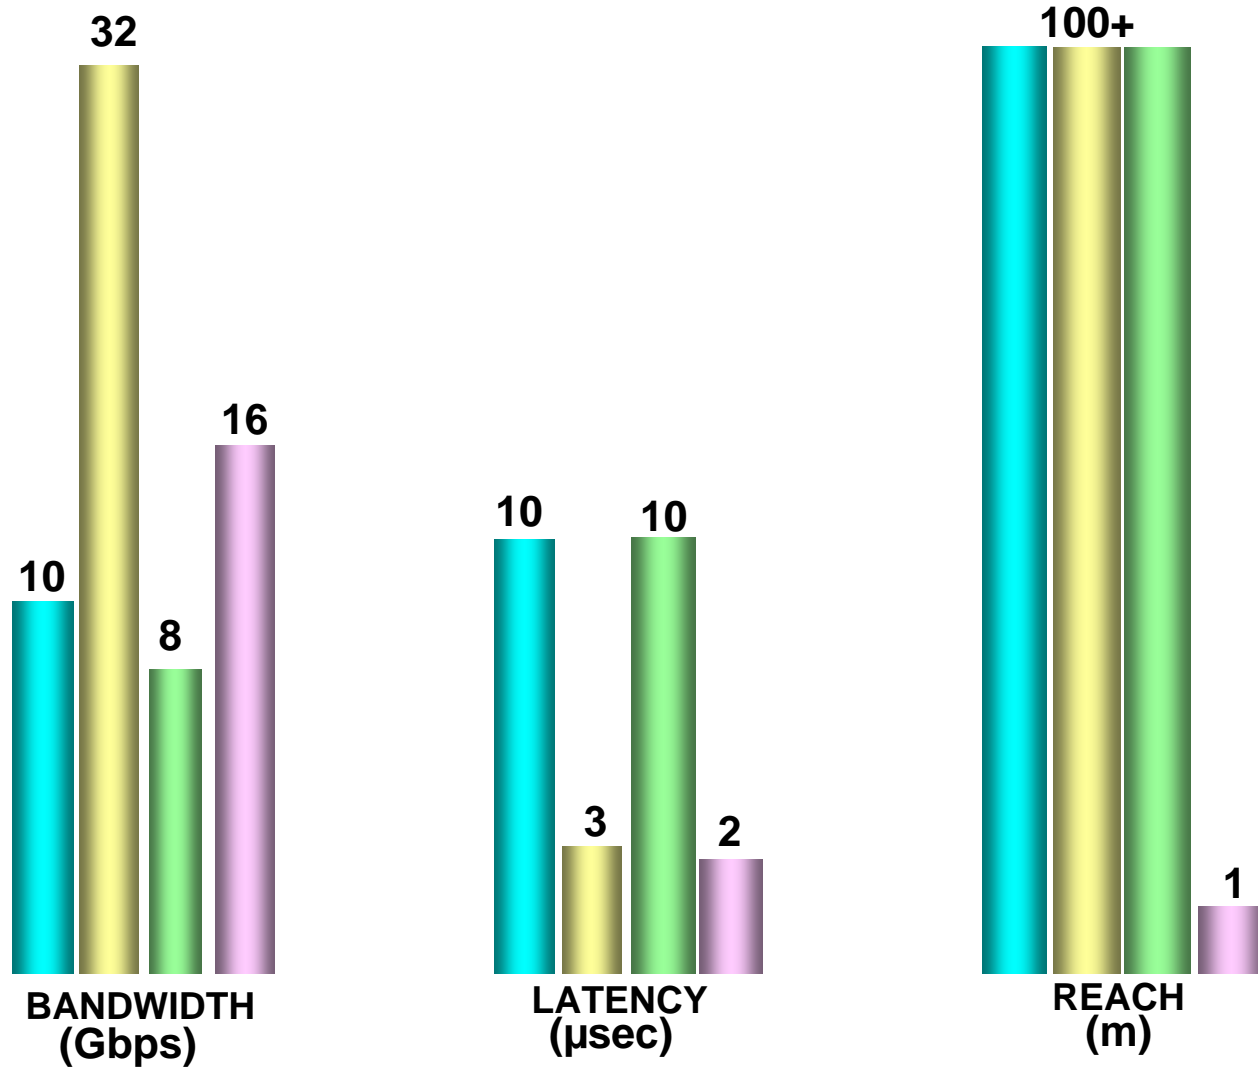
Public Data, Various Sources

Fabric Comparison

10 Gb/s Ethernet may become a 3rd place fabric (for some applications)

- EN, 10Gbps
- IB, 4xQDR
- FC, 8
- PCIe gen2 x4

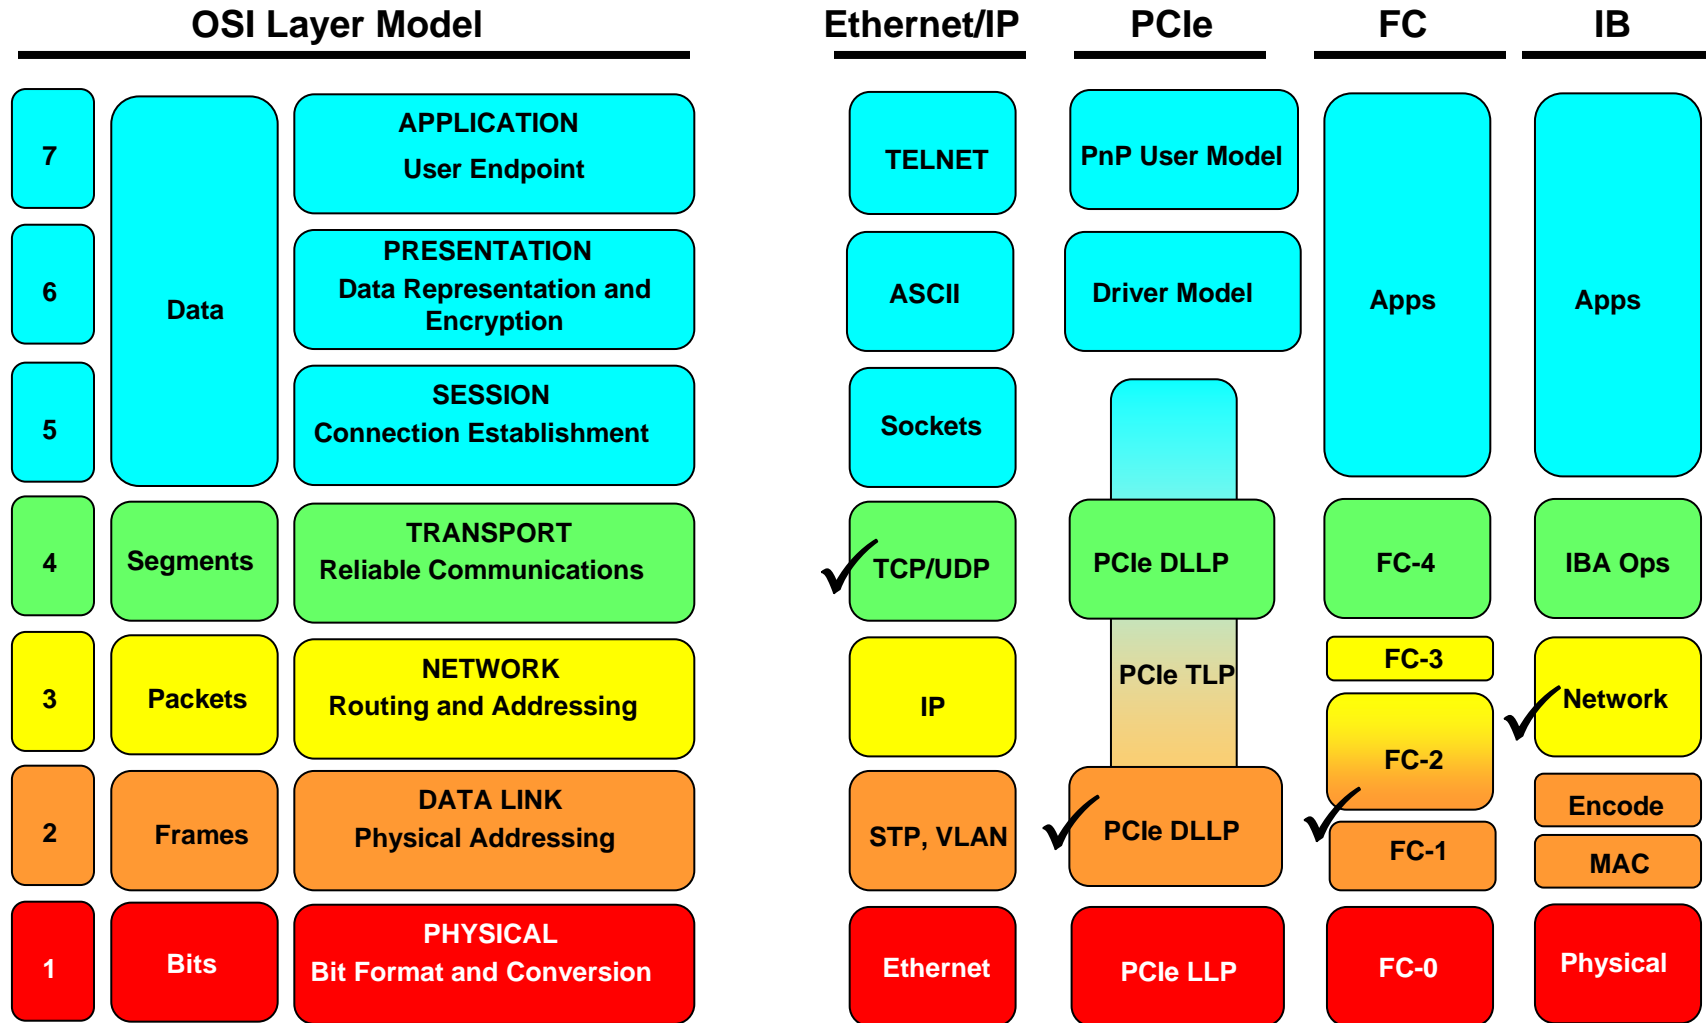
2010
(actual data rates)



Public Data, Various Sources

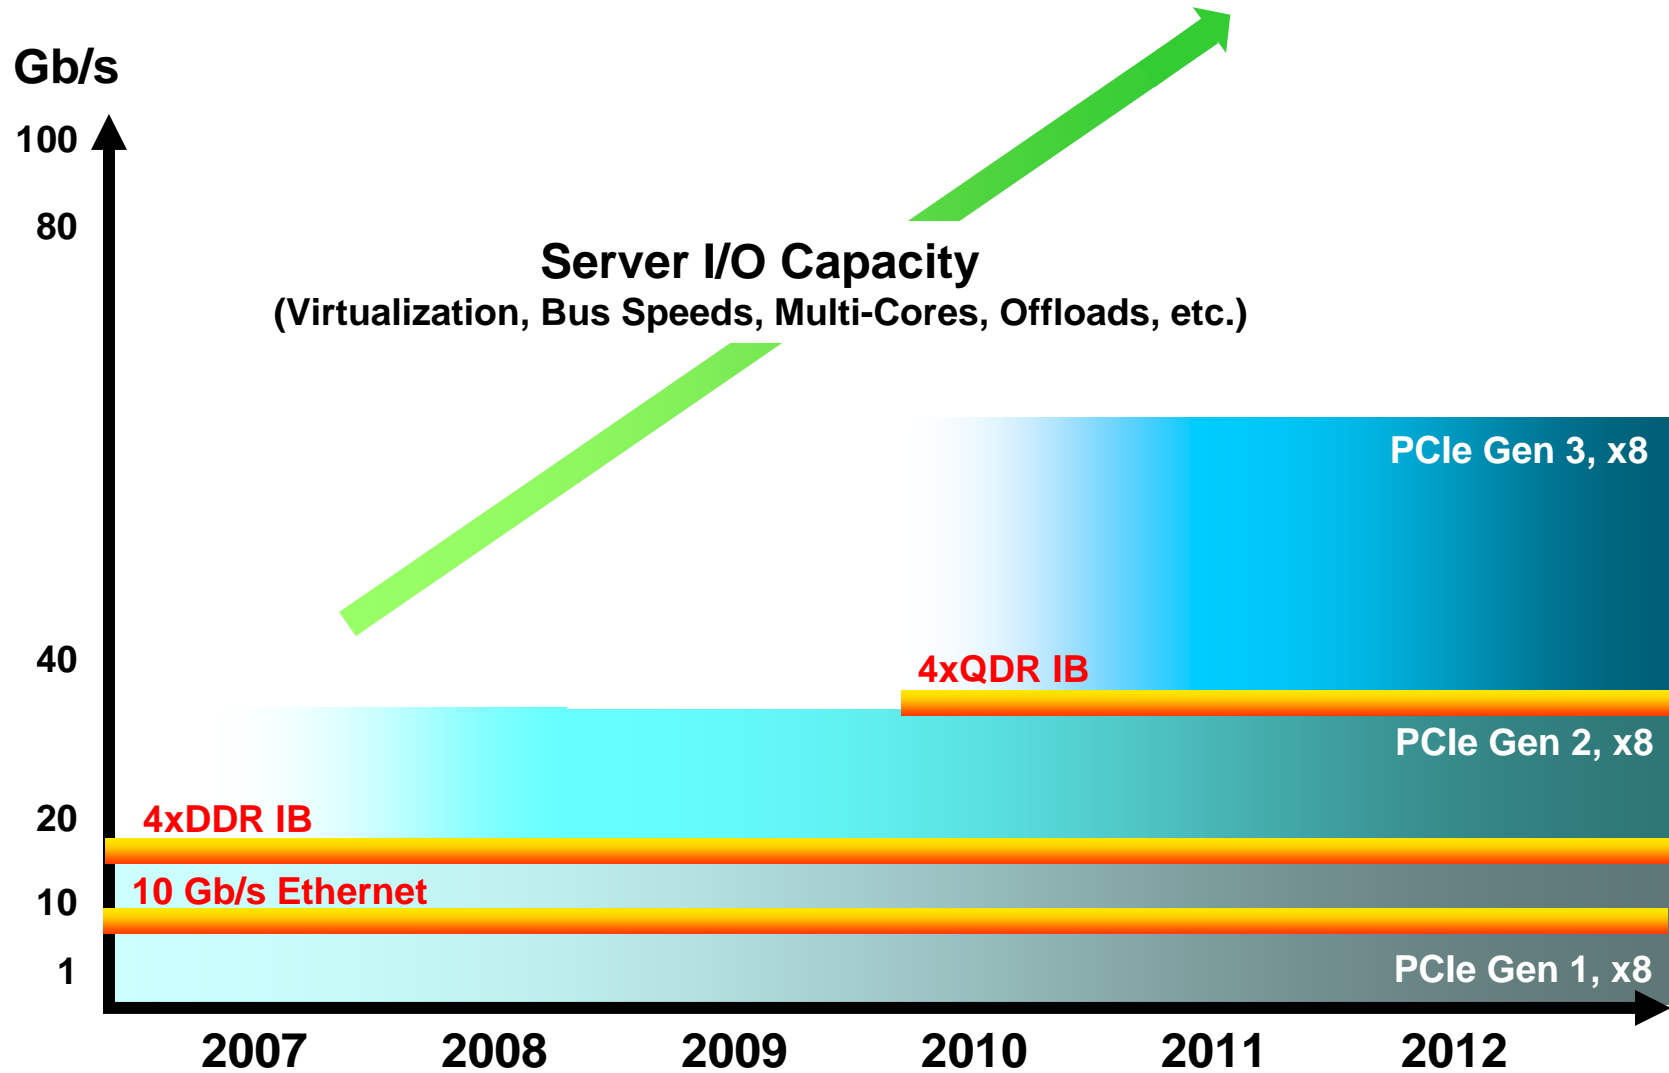
Fabric Comparison

Ethernet is disadvantaged at Layer 2 (compared to other fabrics)



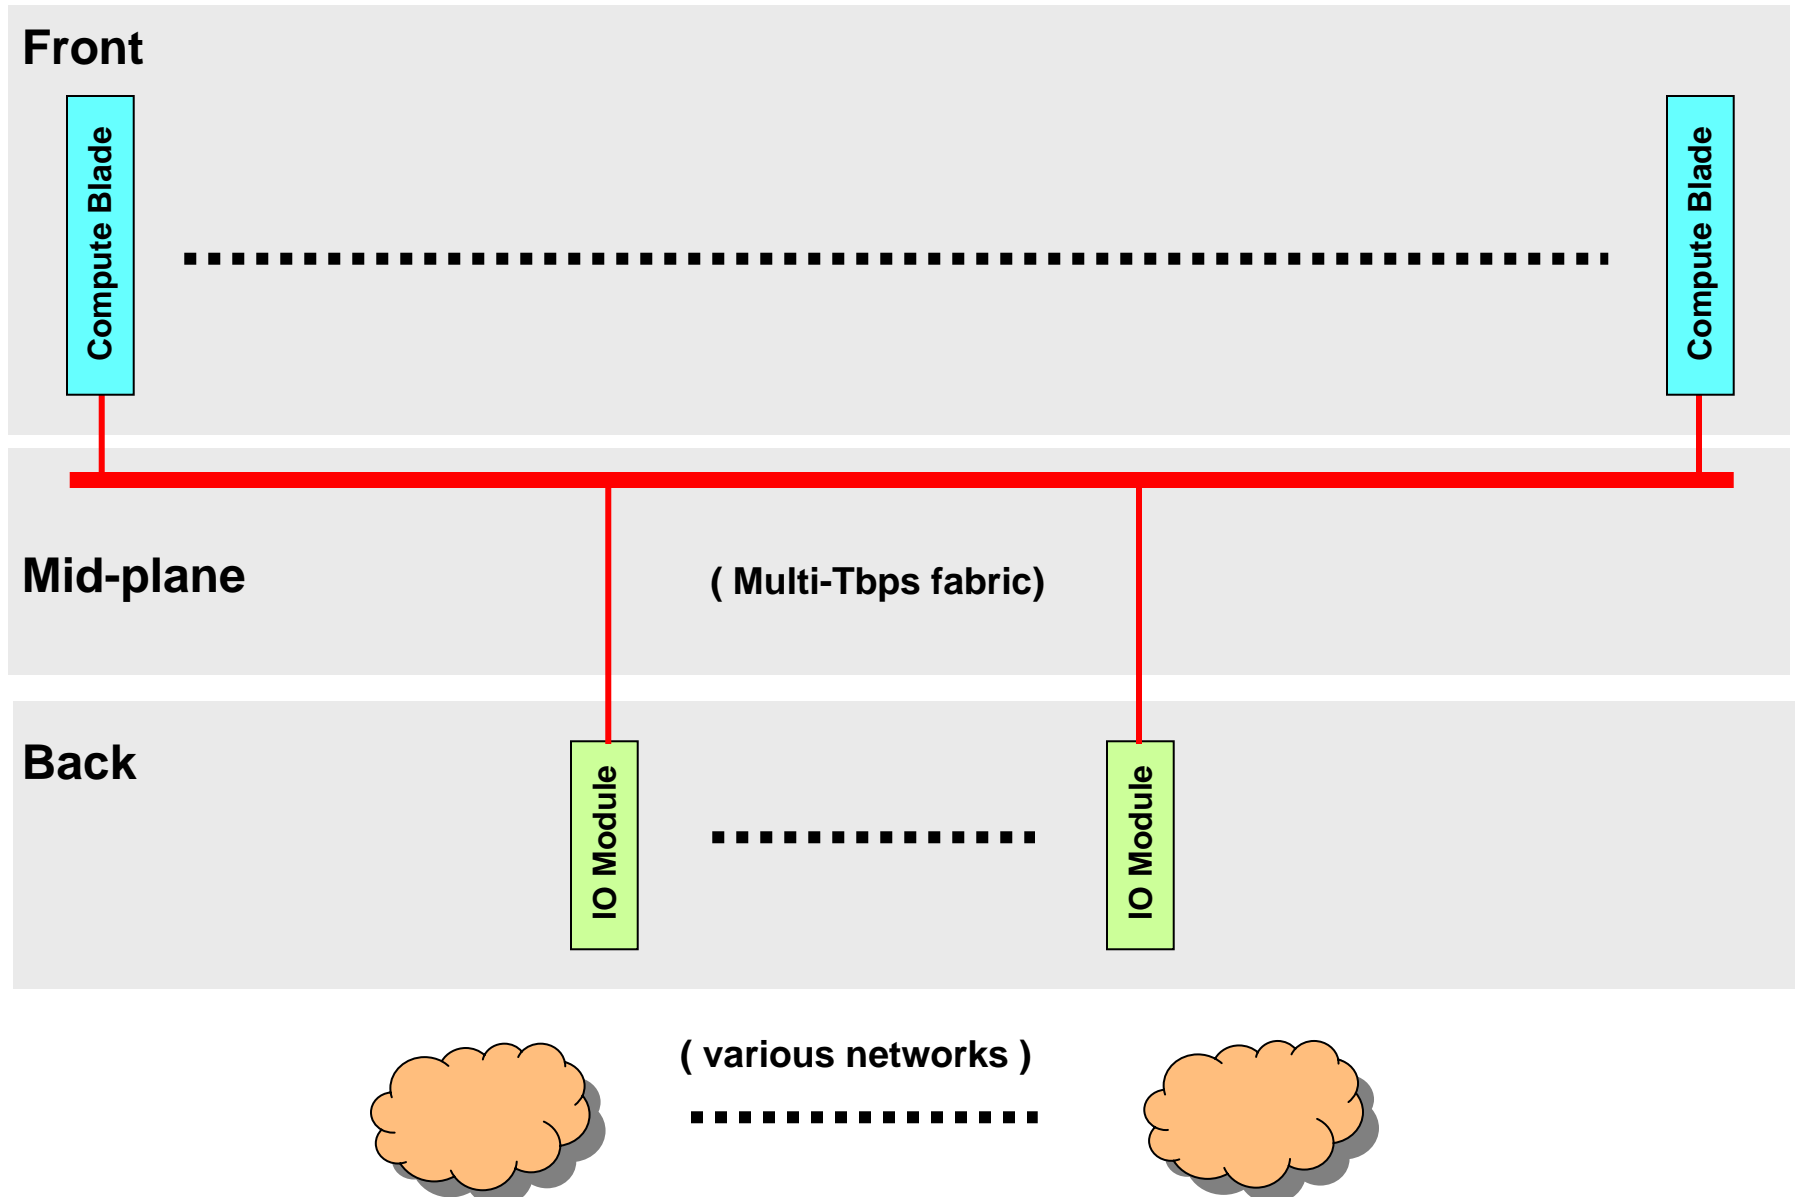
✓ Layer Responsible for Guaranteed Packet Delivery

Server I/O Drivers

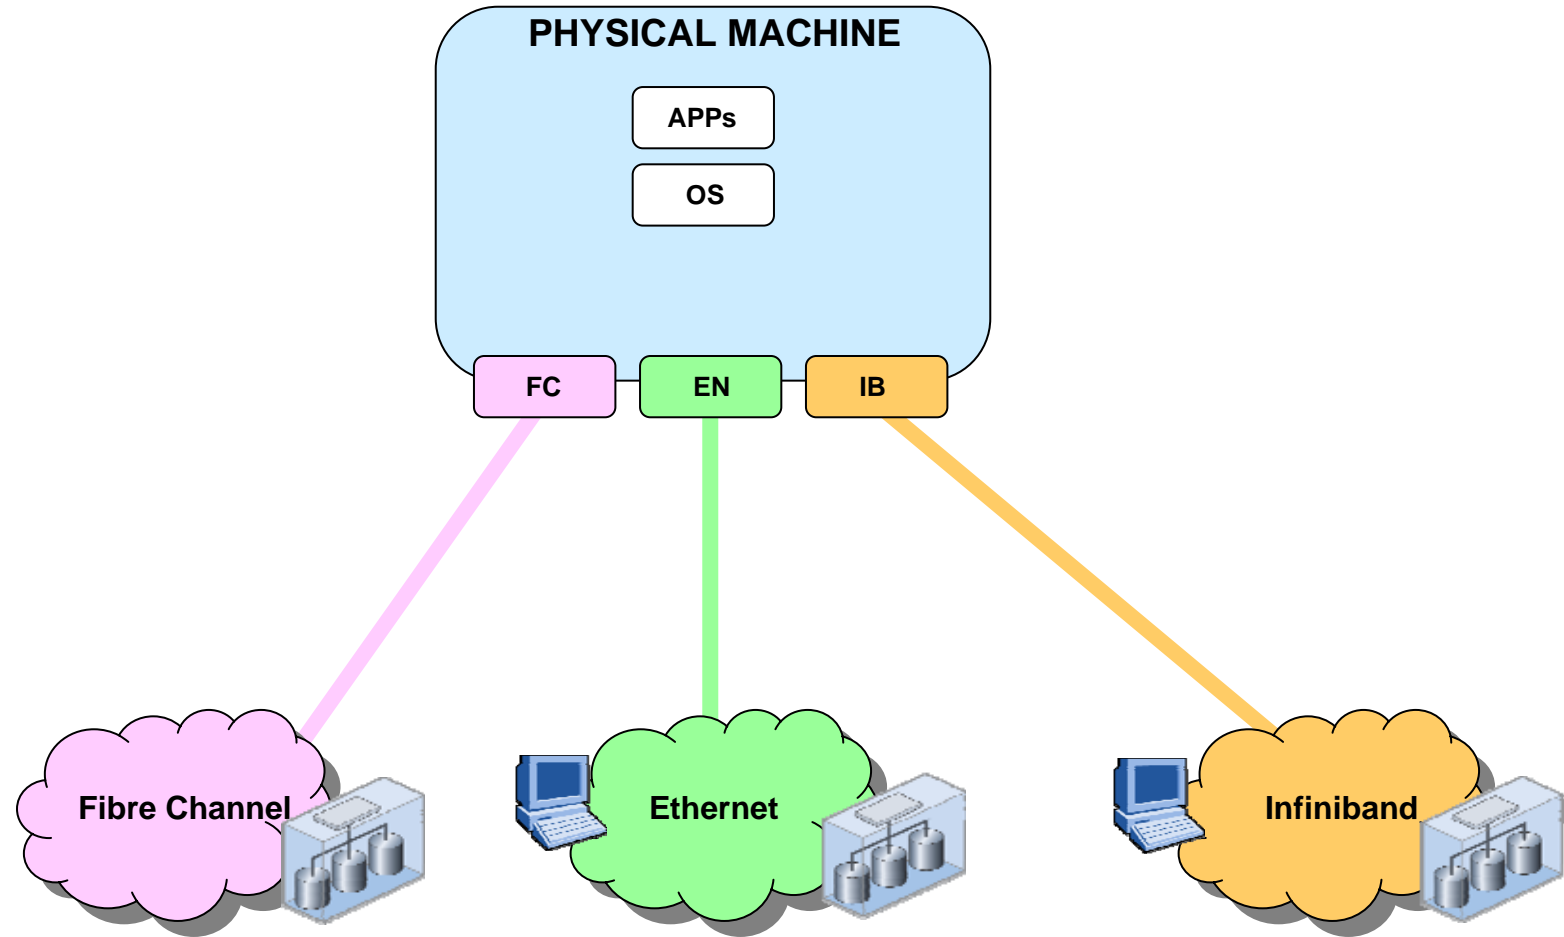


Public Data, Various Sources

General Blade Server Chassis Architecture




Virtualization and Ethernet



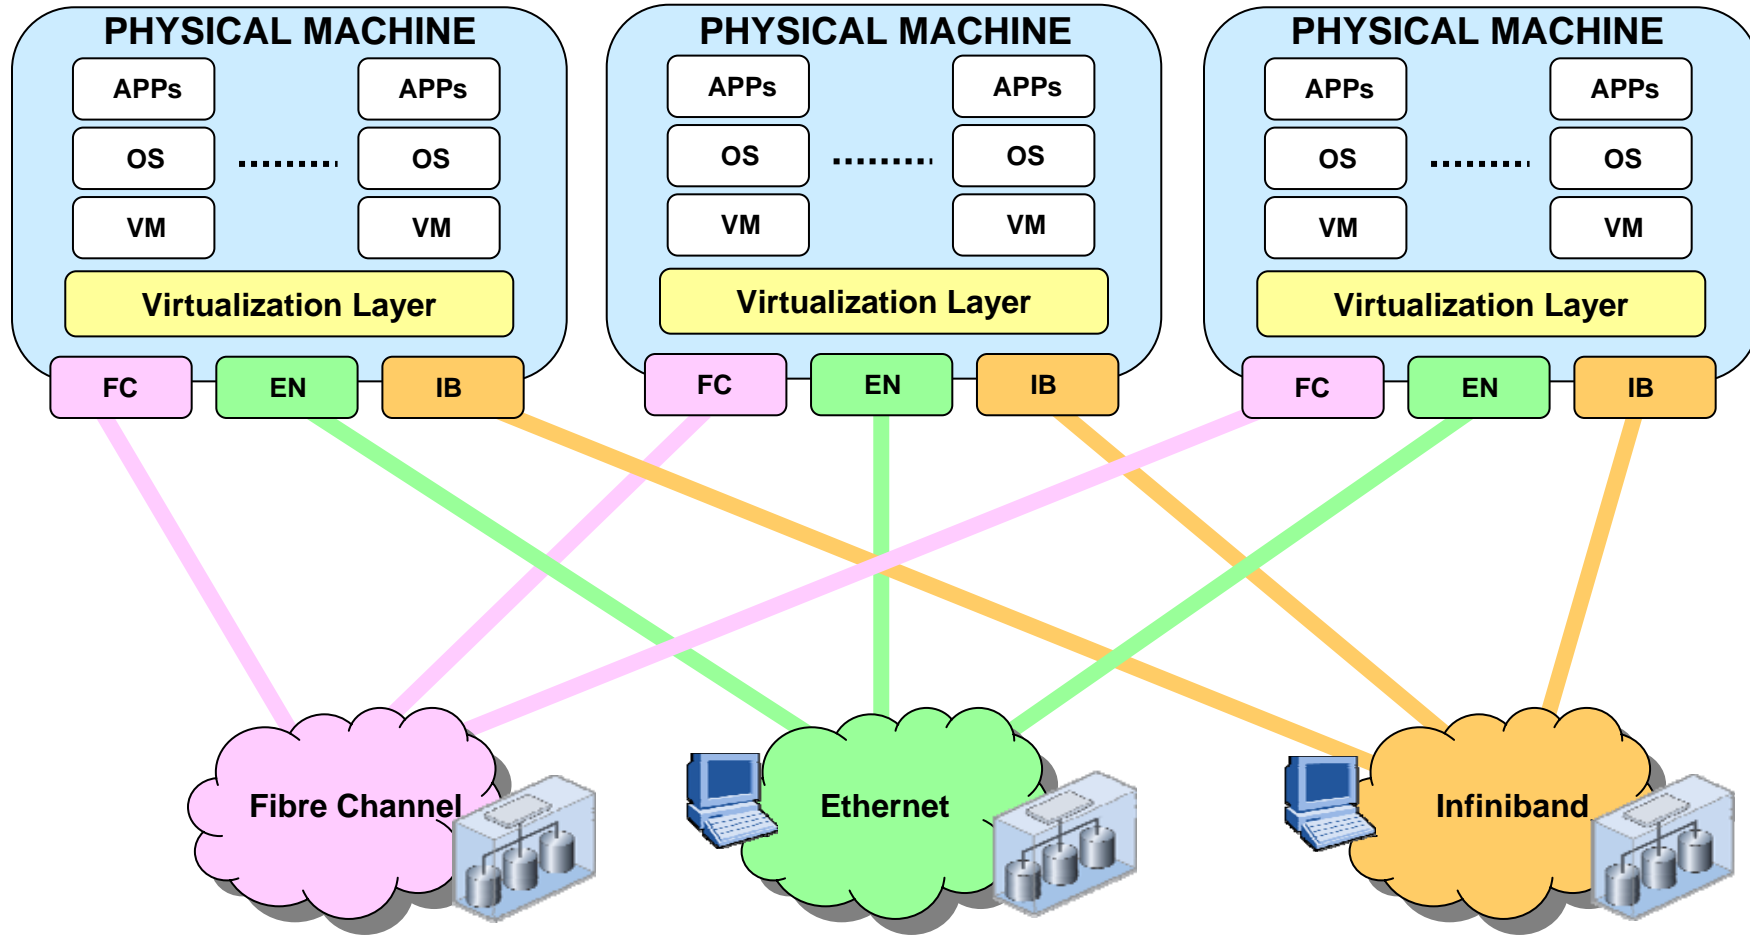
Single OS on a single physical machine
Each type of network access requires a different type of I/O interface

Virtualization and Ethernet



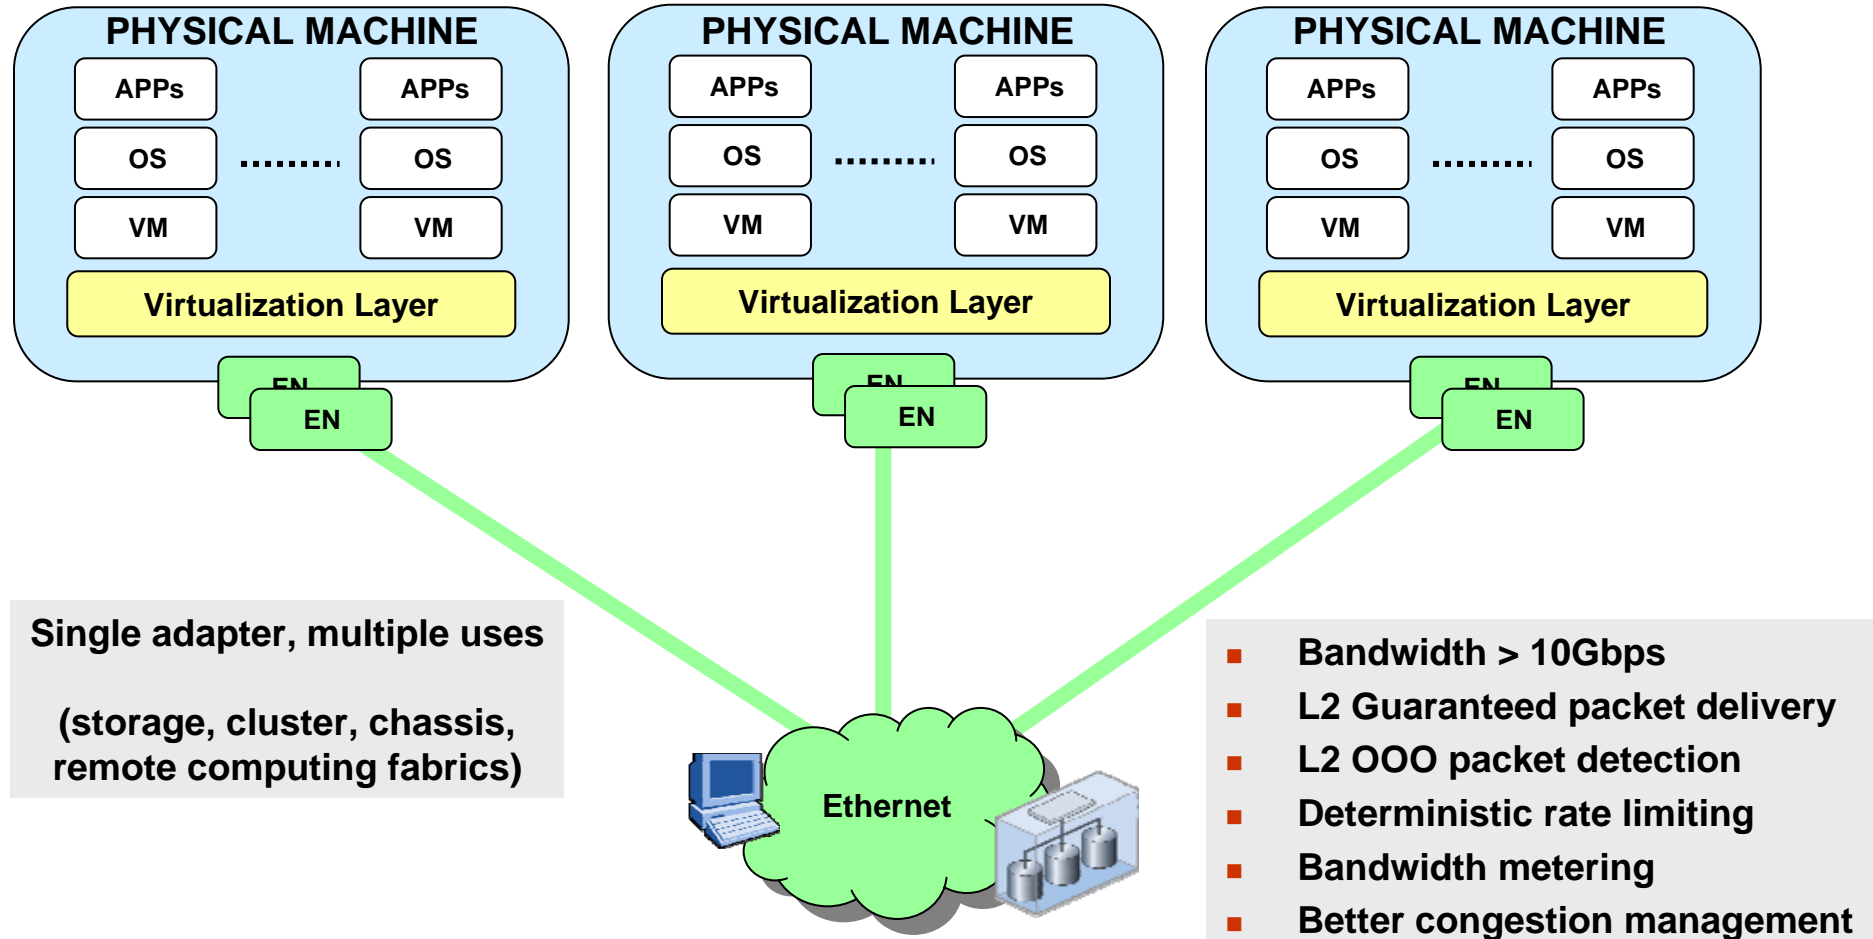
Multiple Oses on a single physical machine
Each OS has a virtual “copy” of each I/O interface
Each type of network access requires a different type of I/O interface

Virtualization and Ethernet



Multiple Oses on multiple physical machines
Each OS has a virtual “copy” of each I/O interface
Each type of network access requires a different type of I/O interface

Virtualization and Ethernet



Multiple OSES on multiple physical machines
Each OS has a virtual “copy” of each I/O interface
Ethernet becomes a “Unifying Fabric”

40GE Drivers

- PCIe Gen 1 x8 consumes 16 Gb/s
- PCIe Gen 2 x8 will not fit in 2 x 10 Gb/s pipe (32 Gb/s)
- PCIe Gen 3 x8 will not fit in 4 x 10 Gb/s pipes (64 Gb/s)
- Chassis will be developed that may have multiple KRAs
 - A single mezzanine card can then push > 10 Gb/s
- Eliminate a major advantage IB has over EN

100GE Drivers

- **Needed for x10 aggregation of 10GE**
 - Else 10GE is limited to stackable/pizza-box switching
- **Needed to push Ethernet into the Telecom space**
 - Achievable “Ethernet Everywhere” Internet

40GE/100GE Co-Existence

- **Dell sees a need for both 40 Gb/s and 100 Gb/s**
 - **If achieving one doesn't seriously impact delivery of the other**
 - **As long as we prevent creation of two different Ethernets**
 - **40 Gb/s**
 - **Will happen either inside or outside IEEE (prefer inside)**
 - **Mostly applies to servers and chassis mid-plane**
 - **Will guaranteed L2 delivery be a requirement?**
 - **100 Gb/s**
 - **Mostly applies to 10 Gb/s aggregation and telecommunications**
 - **Credit-based congestion management is primary here**
 - **Some switching issues**
 - **40 Gb/s does not aggregate well into 100 G/bs**
 - **Switch silicon may be needed that supports 10/40/100**