Before the lecture

- Advanced topics in computer networking
- Assignment
- Question-based Teaching Strategy
- Aadditional resources:

http://www.cs.otago.ac.nz/staffpriv/yawen/402.html

Outline

- This Lecture
  - Cloud Computing
  - Data Center Networking

Cloud Computing

- What is Cloud Computing?

  — Cloud computing is the delivery of computing services (servers, storage, databases, software and more) over the Internet ("the cloud"). Companies offering these computing services are called cloud providers and typically charge for cloud computing services.

Cloud Computing

- Benefits of cloud computing
  - Cost
  - Performance
  - Reliability

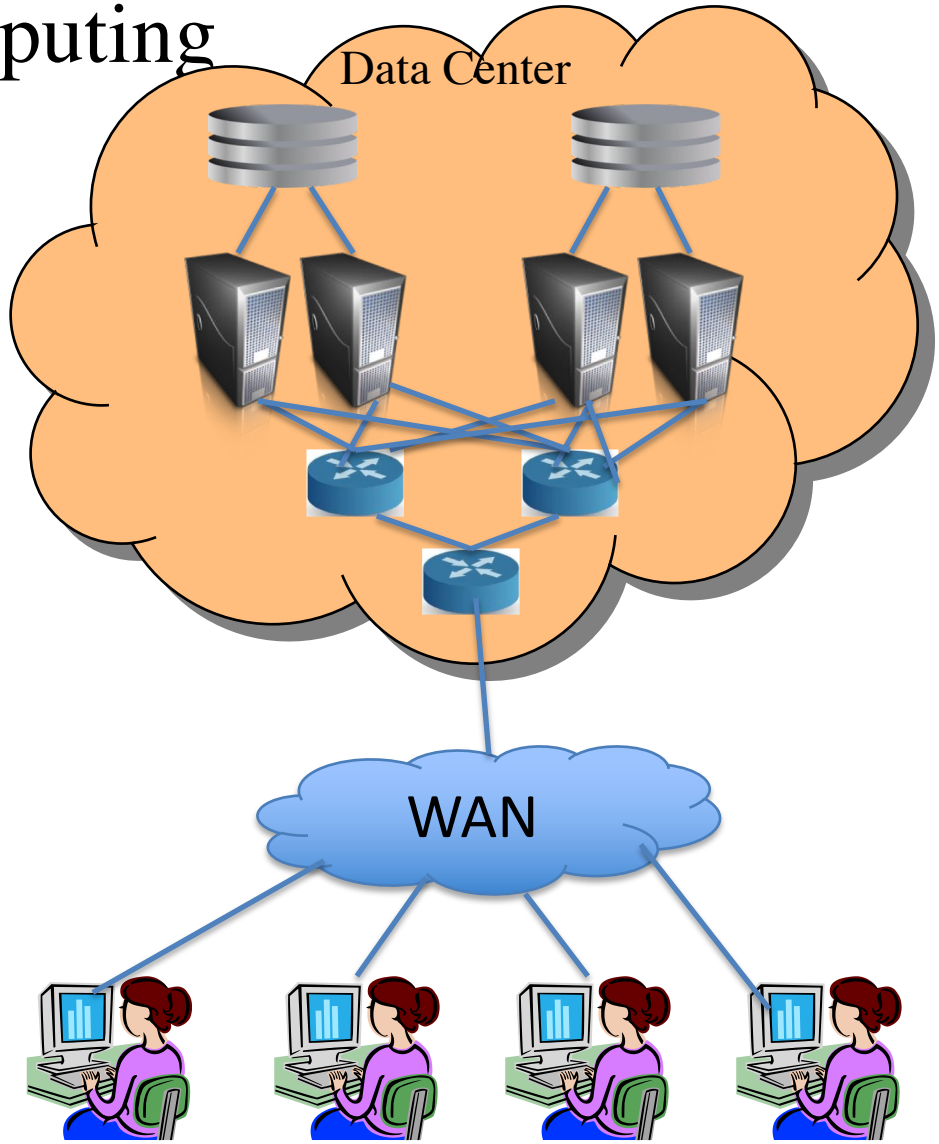

Cloud Computing

- What are SaaS, PaaS, IaaS?

http://www.silverlighthack.com/post/2011/02/27/IaaS-PaaS-and-SaaS-Terms-Explained-and-Defined.aspx

Cloud Computing

- Examples:
  - SaaS: Google Apps

    https://www.google.com/work/apps/business/pricing.html
  - PaaS: Google App Engine

    https://cloud.google.com/appengine/
  - IaaS: Google Compute Engine

    https://cloud.google.com/compute/

Cloud Computing

- Types of cloud deployments: public, private, hybrid
  - Public cloud (owned and operated cloud server provider, which deliver their computing resources like servers and storage over the Internet. You access these services and manage your account using a web browser.)
  - Private cloud (used exclusively by a single business or organization.)
  - Hybrid cloud (combine public and private clouds, bound together by technology that allows data and applications to be shared between them, greater flexibility and more deployment options)

Cloud Computing

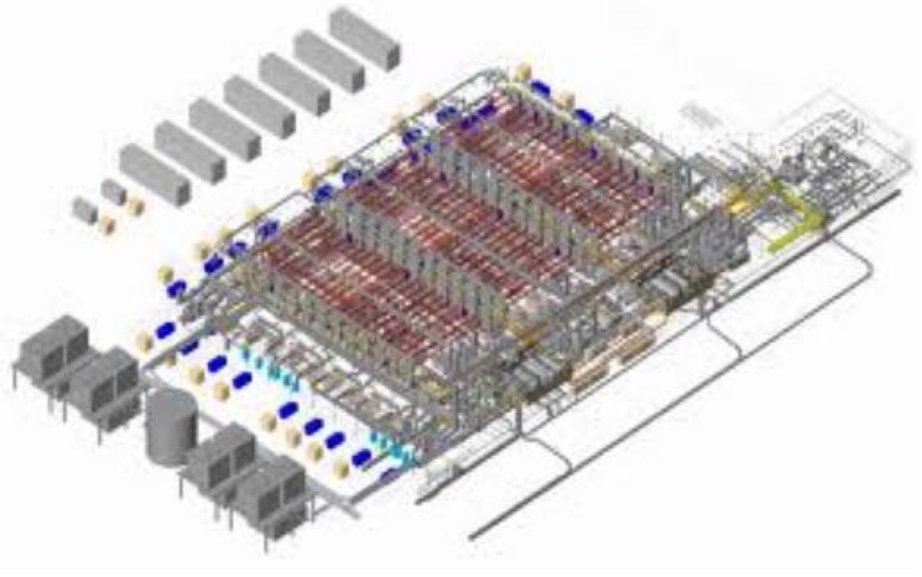

How to connect the computing and storage devices to maximize the performance?

Outline

- This Lecture
  - Cloud Computing
  - Data Center Networking
    - Data Center Physical Layout
    - Main Components in the Data Center
    - Data Center Network Topologies

Data Center Physical Layout

Google's Data Center

Cooling Plant

Data Center Physical Layout

Data Center Physical Layout

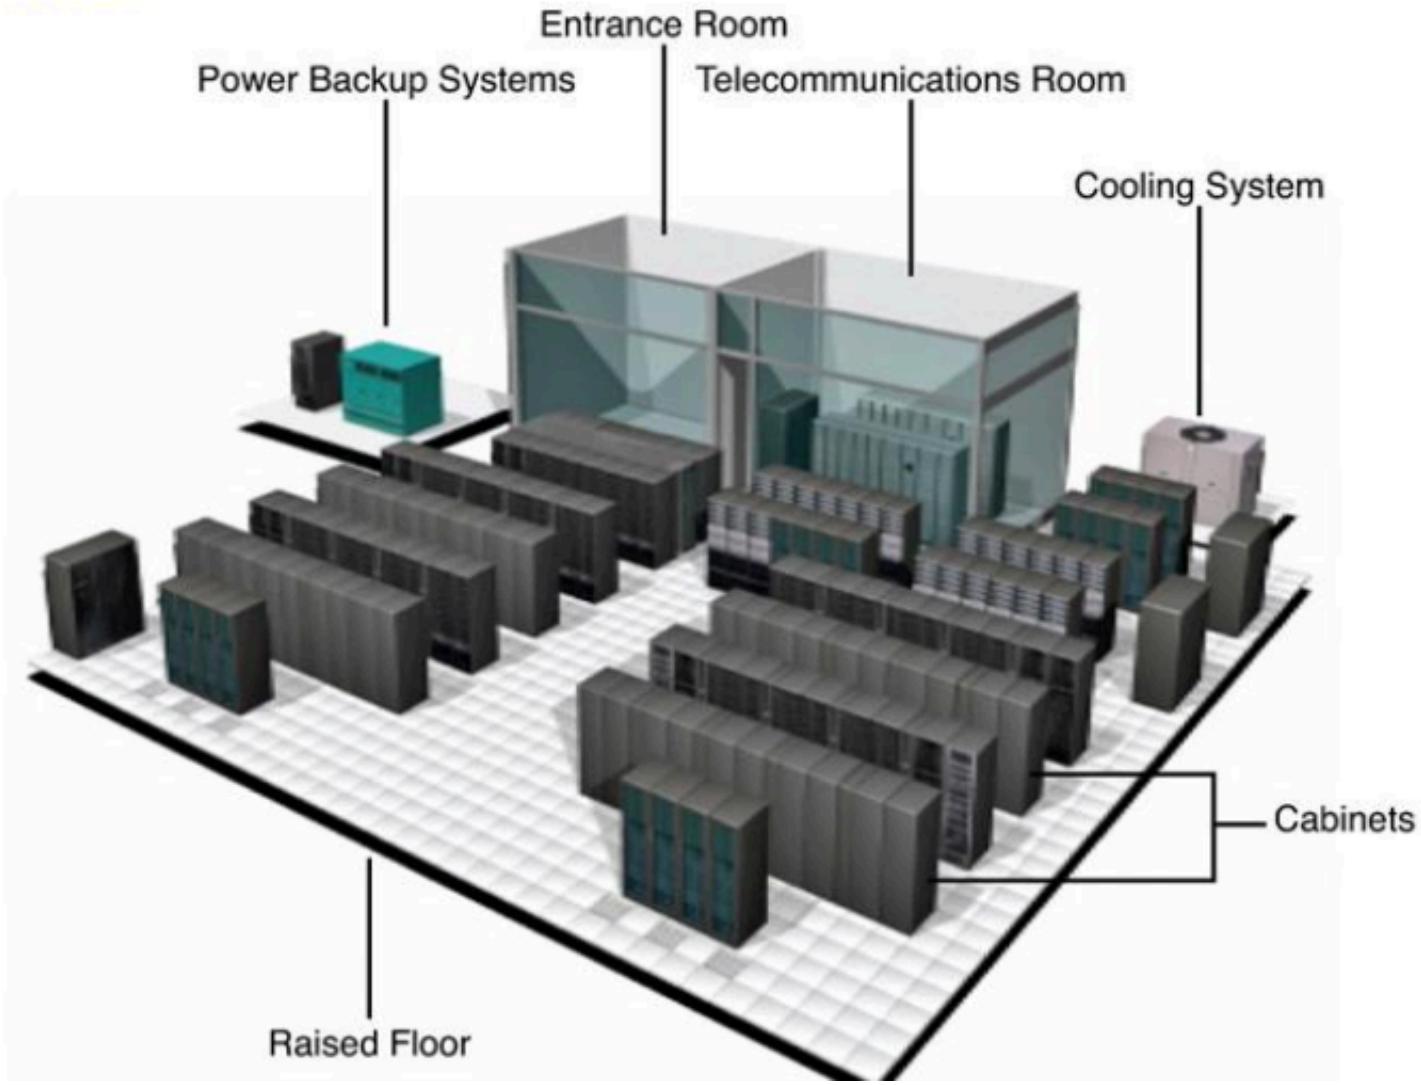

Data Center Physical Layout

- Google released a gallery of pictures of their Data Centers around the world (7 in the Americas, 3 in Asia and 3 in Europe). http://webodysseum.com/technologyscience/visit-the-googles-data-centers/#sthash.zzWAk3DV.dpuf

- Google also published a guided tour on YouTube: https://www.youtube.com/watch?v=avP5d16wEp0

- Street View tour of their data center in Lenoir, North Carolina: http://www.google.com/about/datacenters/inside/streetview/

Outline

- This Lecture
  - Cloud Computing
  - Data Center Networking
    - Data Center Physical Layout
    - Main Components in the Data Center
    - Data Center Network Topologies

Main Components in the Data Center

- **Servers**
  - rack-mount servers
  - blade servers

http://hacksnpasses.blogspot.co.nz/2011/07/pros-and-cons-of-tower-rack-and-blade.html

Main Components in the Data Center

- **Storage**
  - Direct Attached Storage (DAS)
  - Network Attached Storage (NAS)
  - Storage Area Network (SAN)

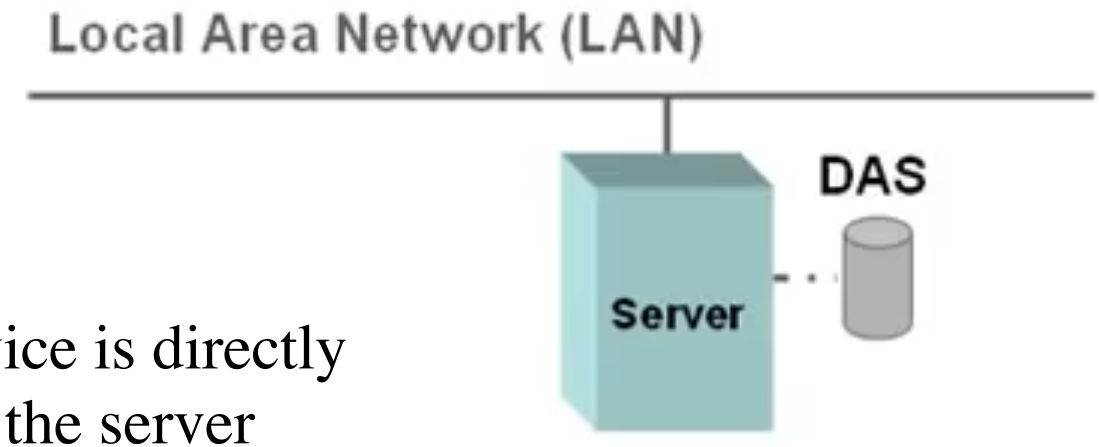

storage device is directly attached to the server

Main Components in the Data Center

- **Storage**
  - Direct Attached Storage (DAS)
  - Storage Area Network (SAN)
  - Network Attached Storage (NAS)

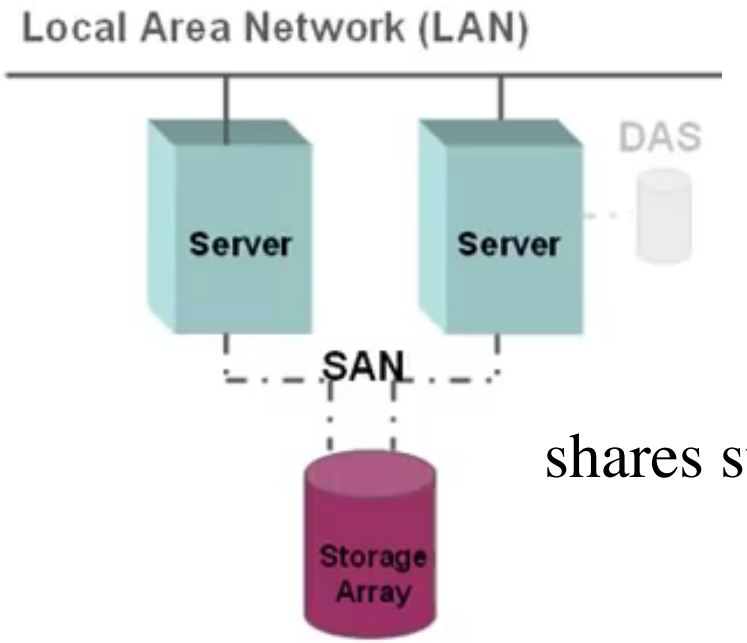

shares storage over a network.

Main Components in the Data Center

- **Storage**
  - Direct Attached Storage (DAS)
  - Storage Area Network (SAN)
  - Network Attached Storage (NAS)

dedicated storage network transfers data between storage devices and servers.

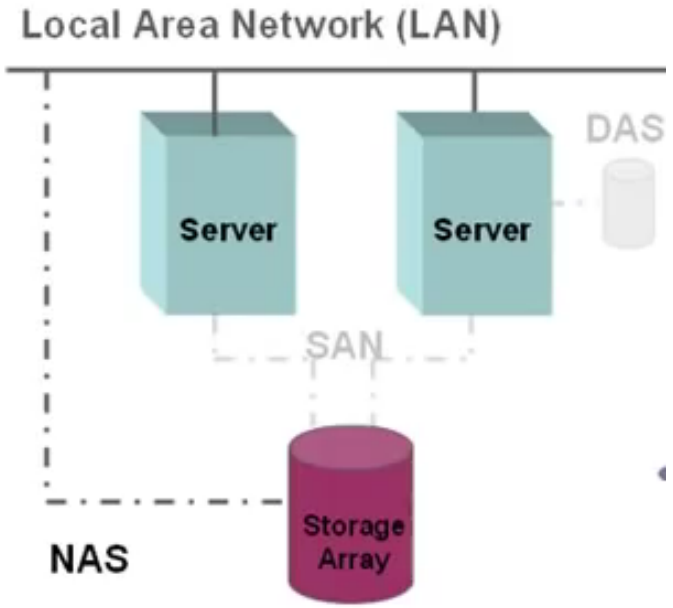

Evolution of Network Storage

Direct Attached Storage

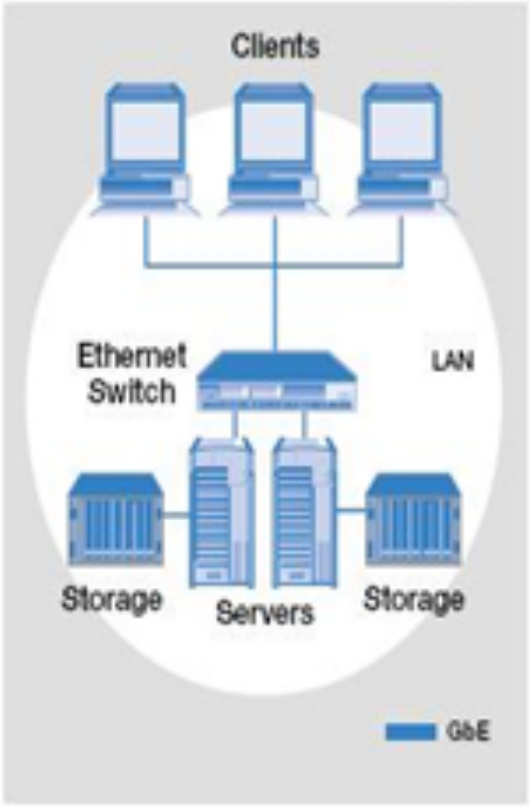

- High cost of management
- Inflexible
- Expensive to scale

Network Attached Storage

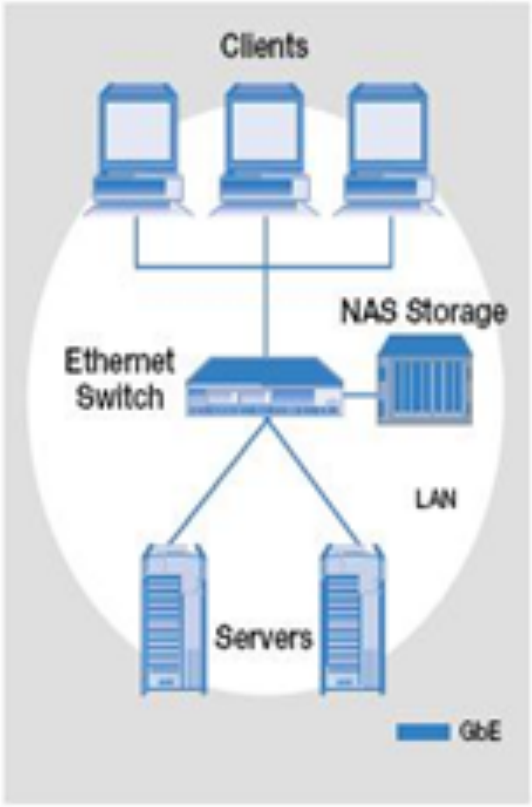

- Transmission optimized for file transactions
- Storage traffic travels across the LAN

Storage Area Network

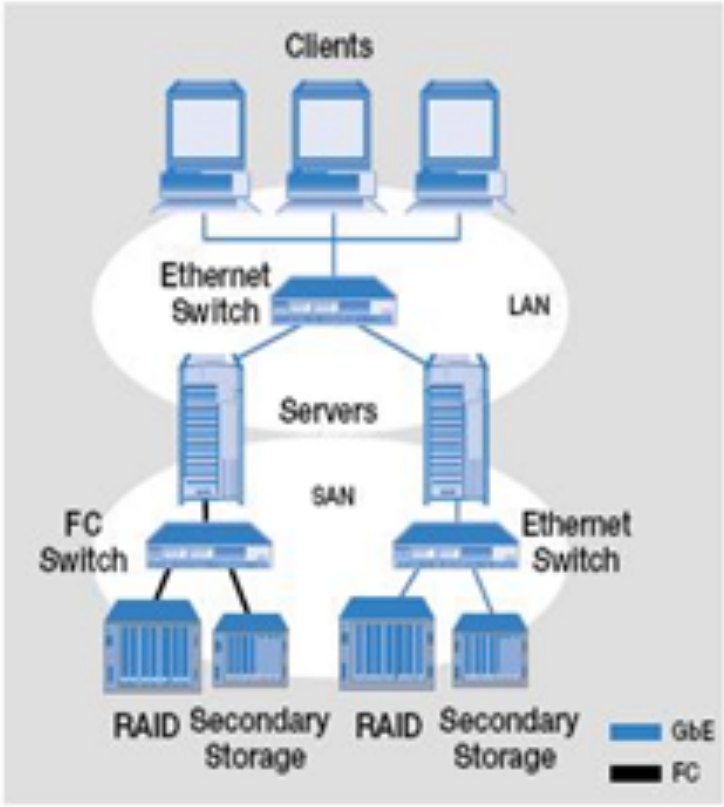

- Transmission optimized for file transactions
- Separate LAN and SAN
- Increases data availability
- Flexible and scalable

Outline

- This Lecture
  - Networking Trends
  - Cloud Computing
  - Data Center Networking
    - Data Center Physical Layout
    - Main Components in the Data Center
    - Data Center Network Design

Data Center Network Design

Racks of servers (Google)

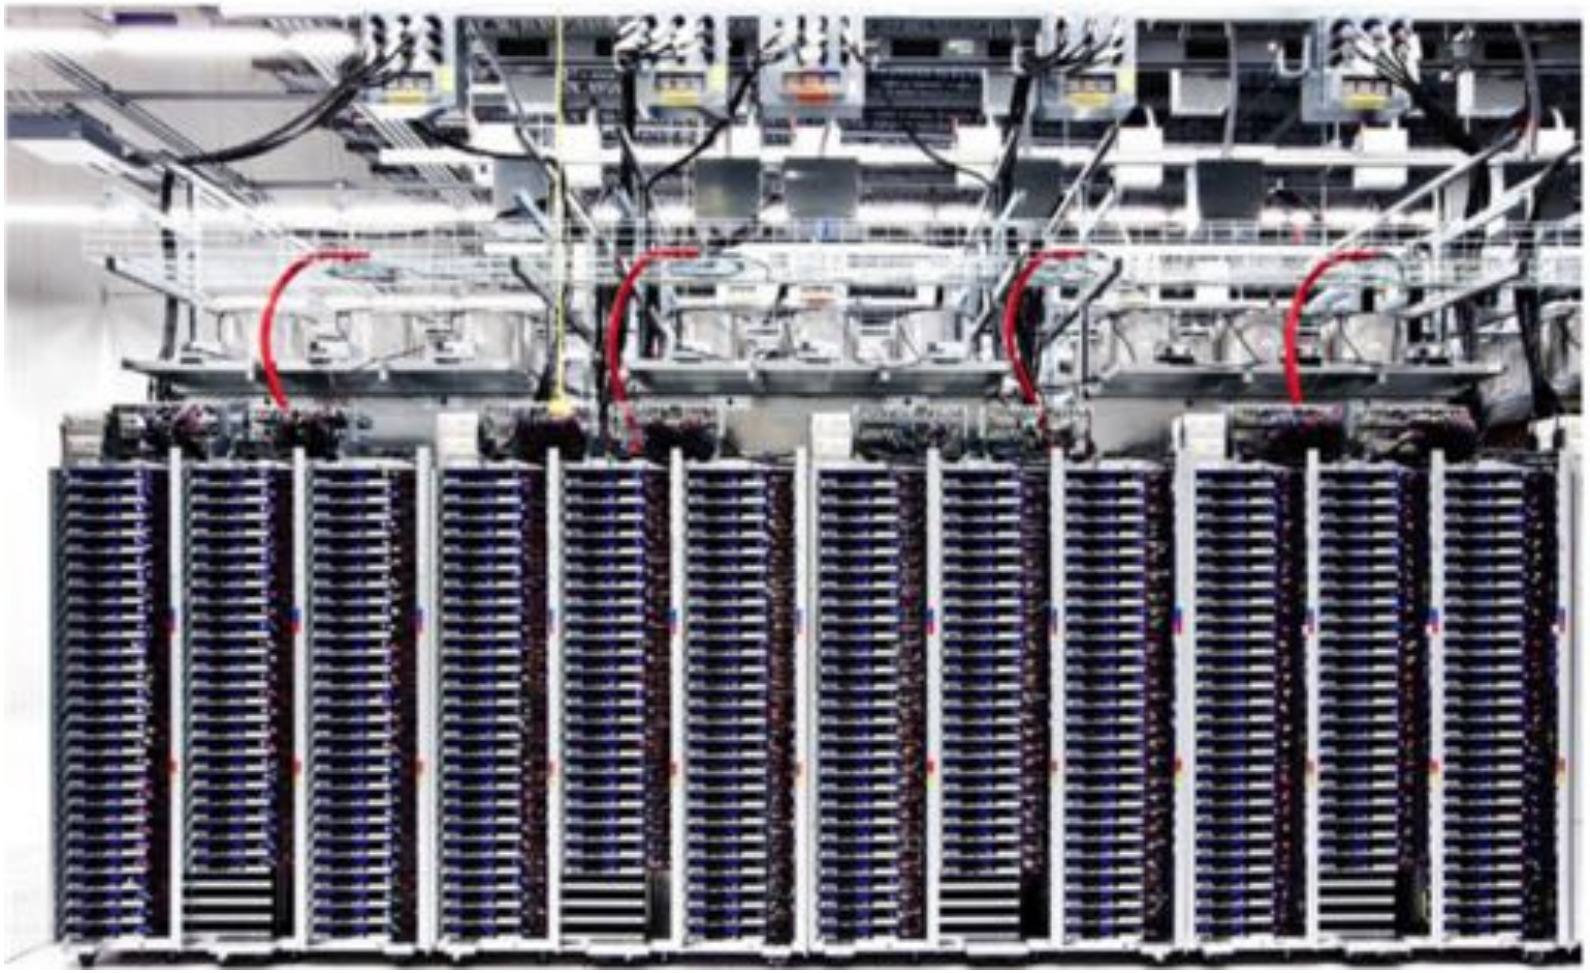

Data Center Network Design

Facebook

What are the connectivity designs?

Data Center Network Design

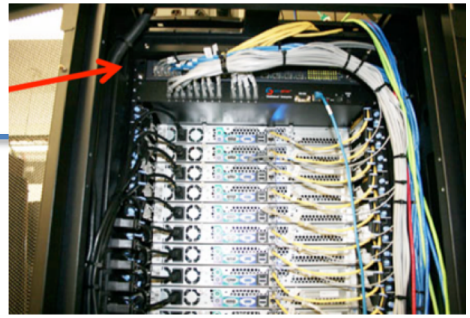

- ToR

TOR – Top of Rack design:

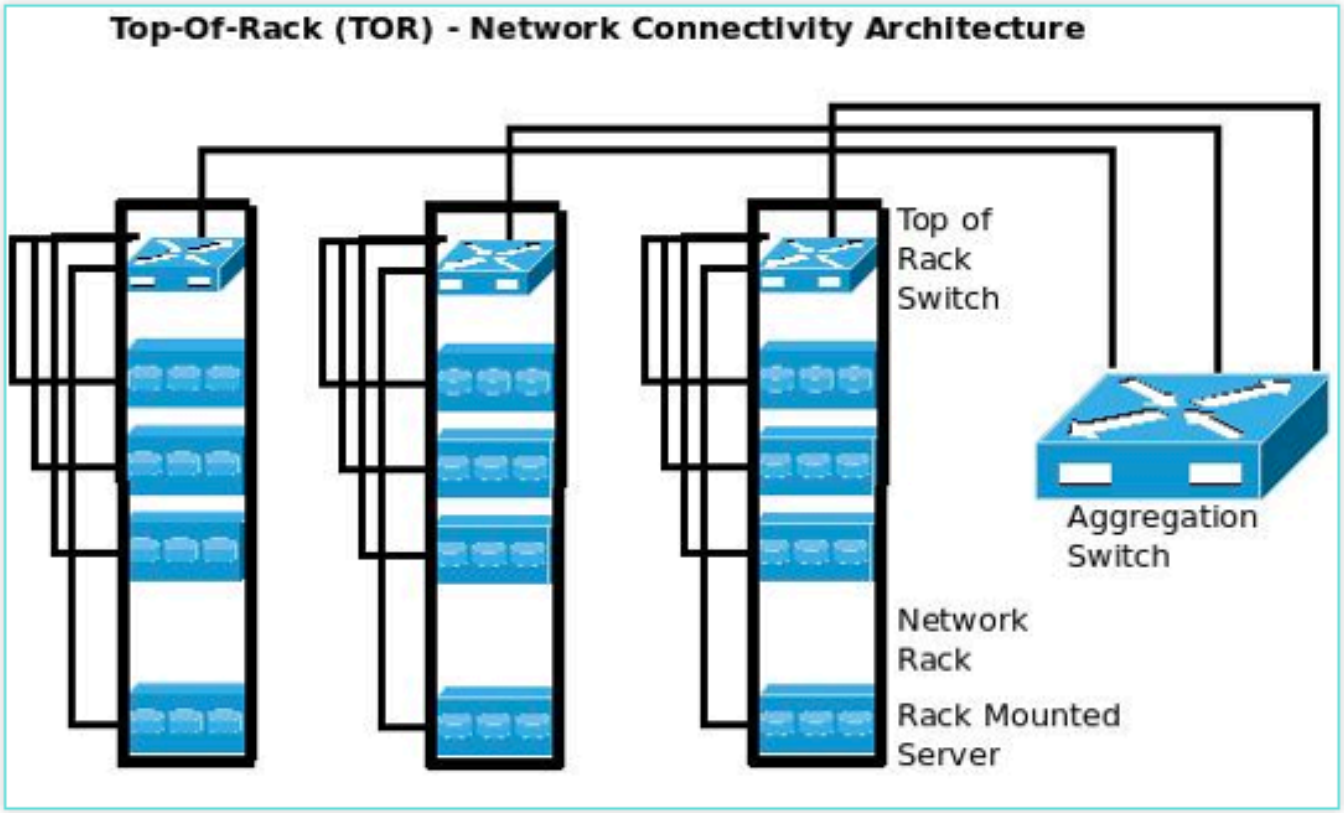

Data Center Network Design

- ToR
  - Easier cabling
  - If rack is not fully populated, unused ToR ports
  - If rack traffic demand is high, difficult to add more ports
  - Upgrading (1G to 10G) requires complete Rack upgrade

Data Center Network Design

- EoR

EOR – End of Row design:

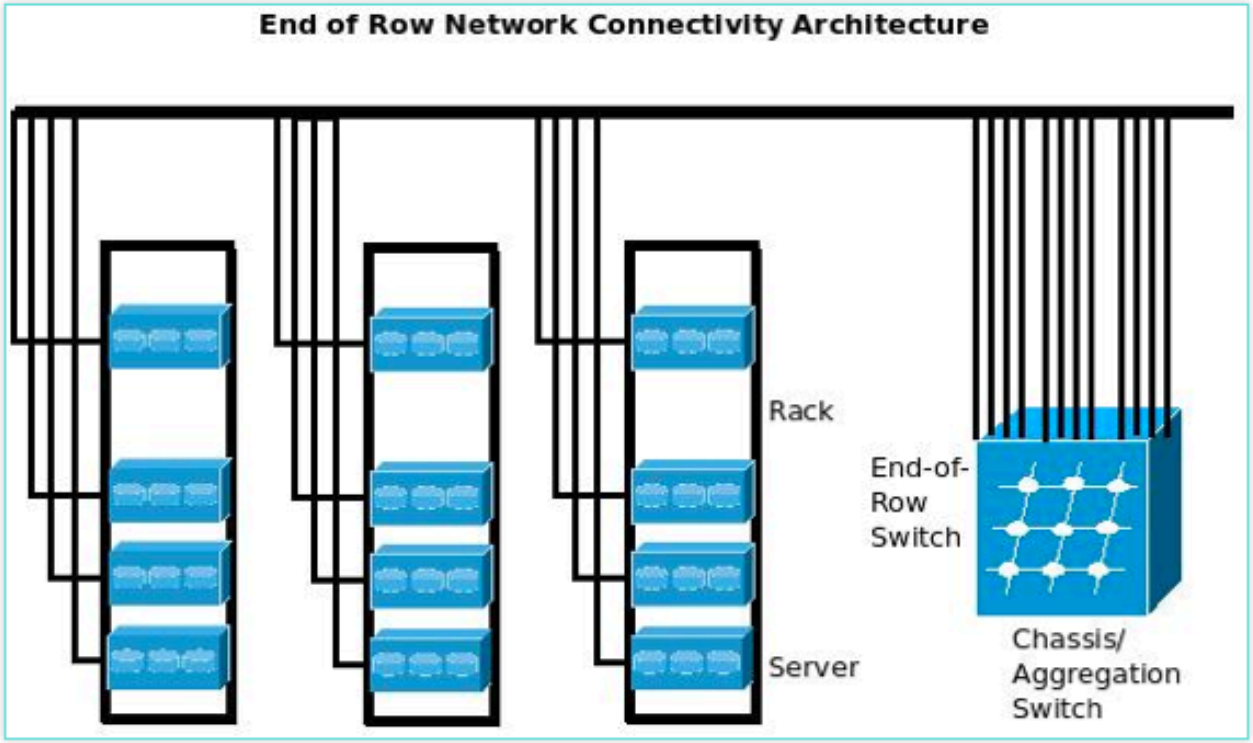

Data Center Network Design

- EoR:
  - Longer cables
  - Severs can be placed in any rack
  - Ports can be easily added, upgraded

Data Center Network Design

- Data Center Design

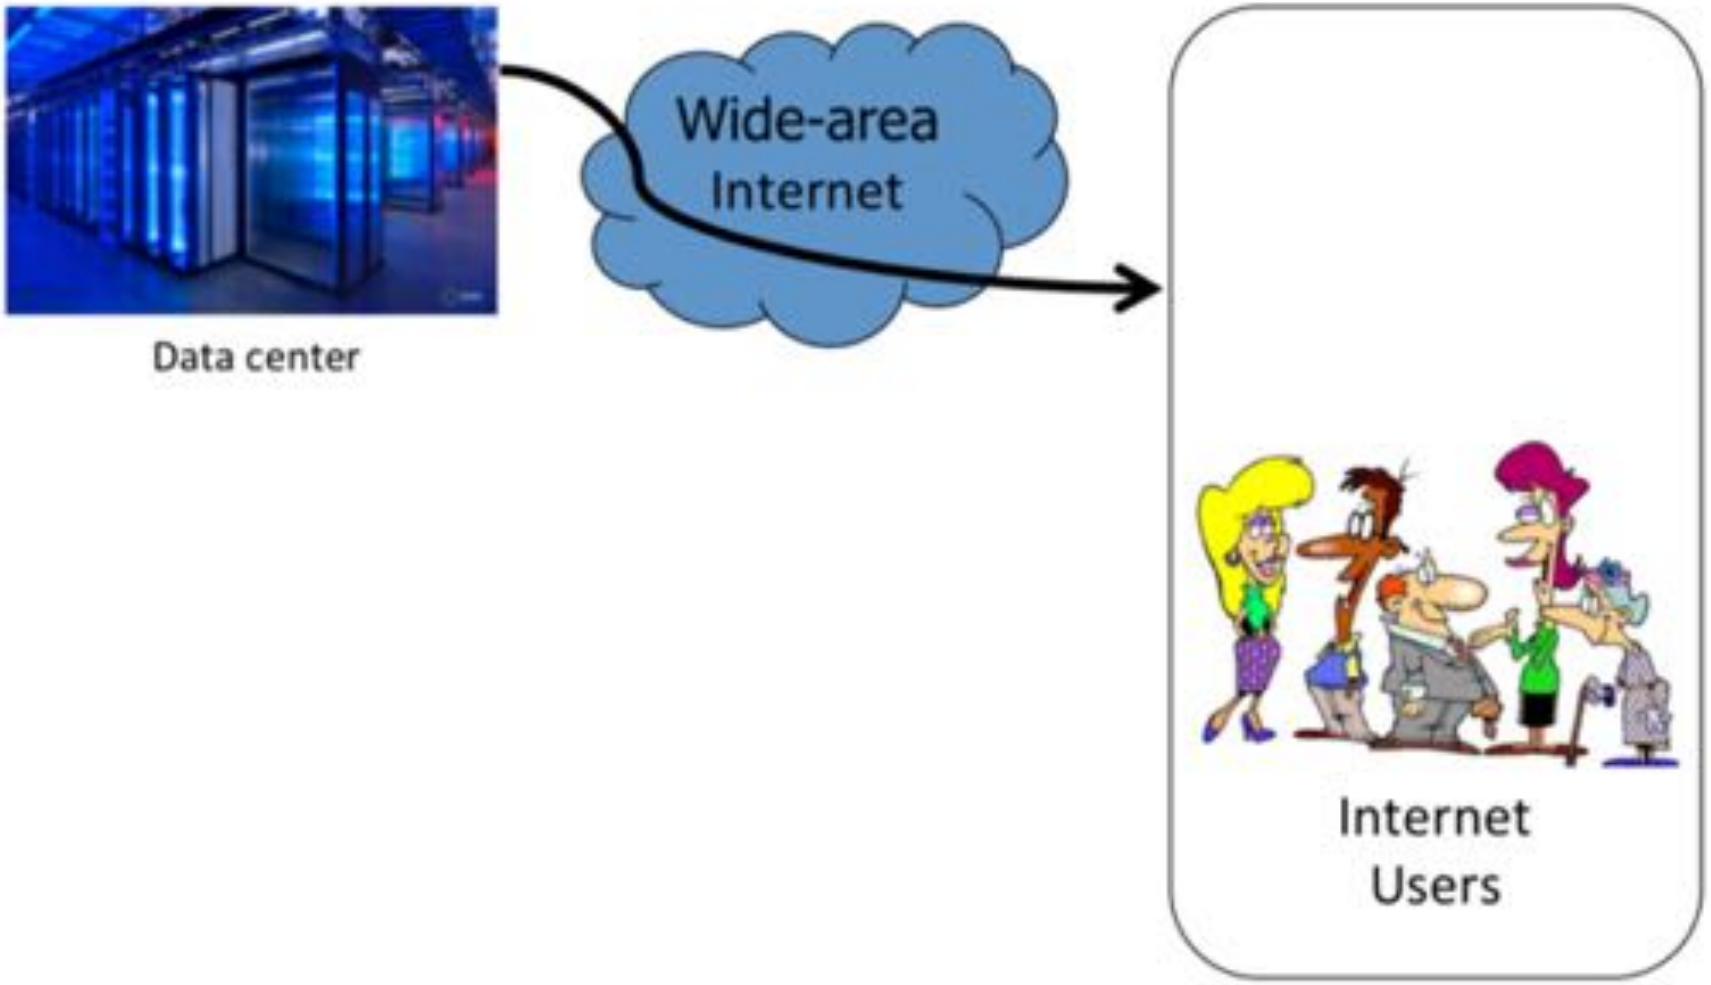

Data Center Network Design

- Data Center Design

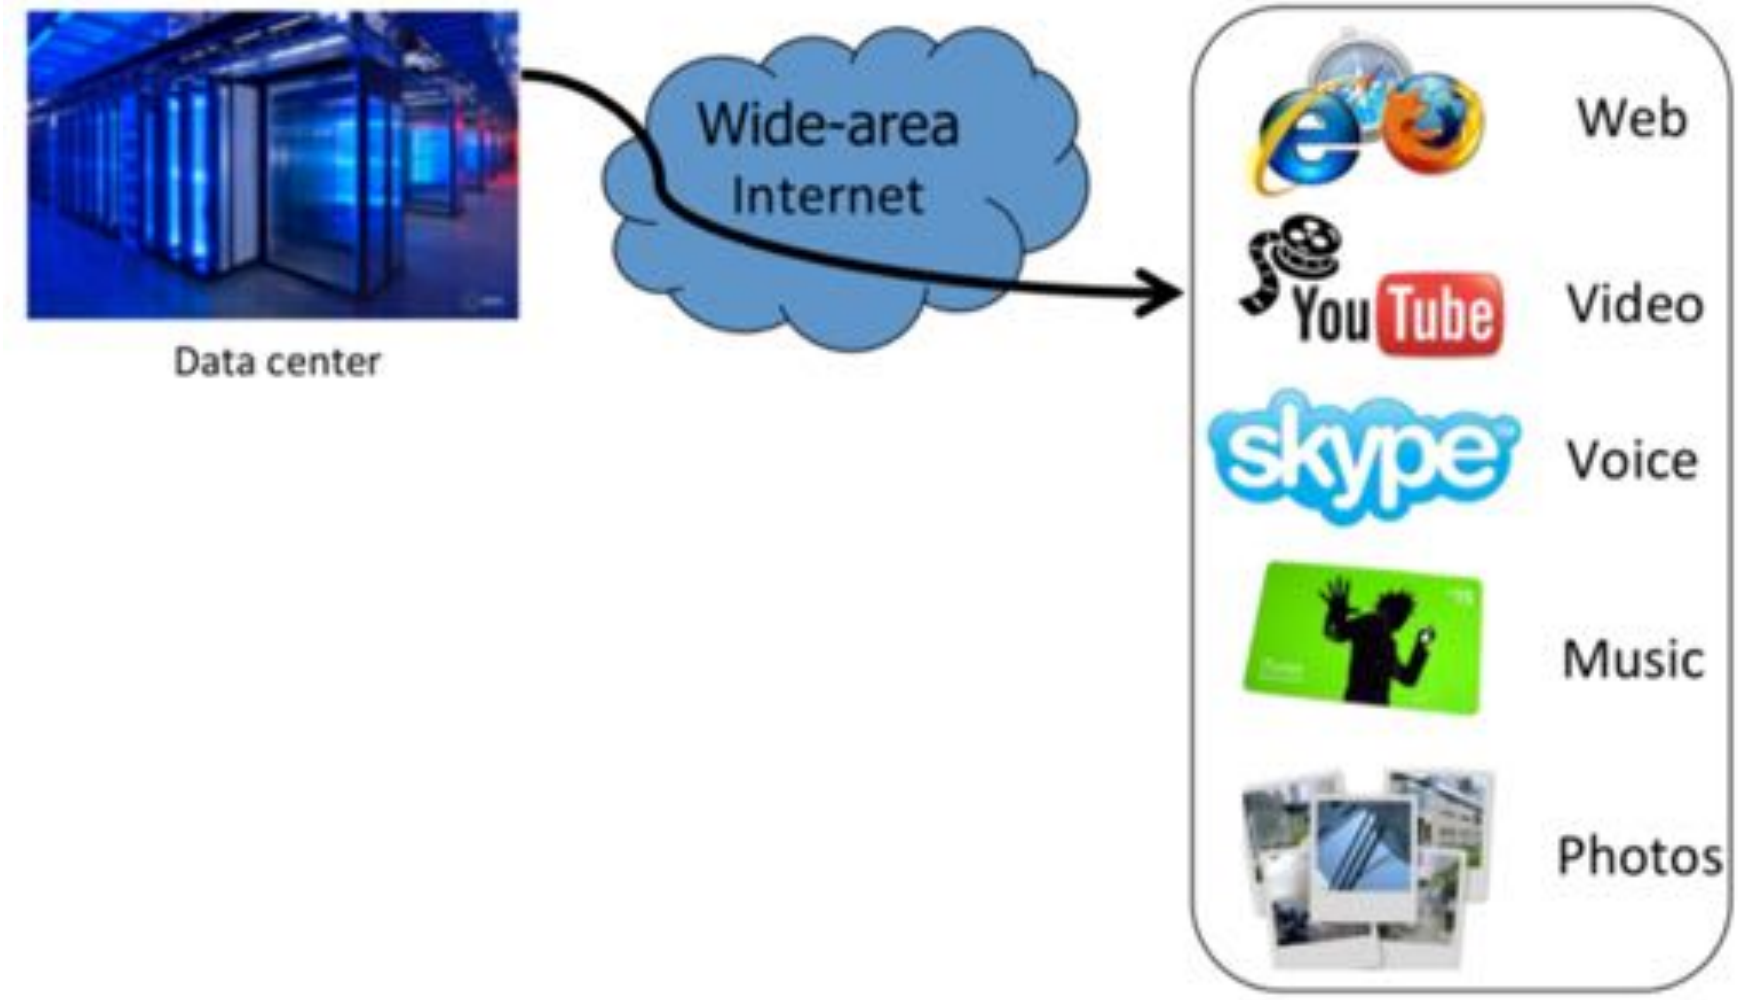

Data Center Network Design

- Multi-Tier Data Center Design

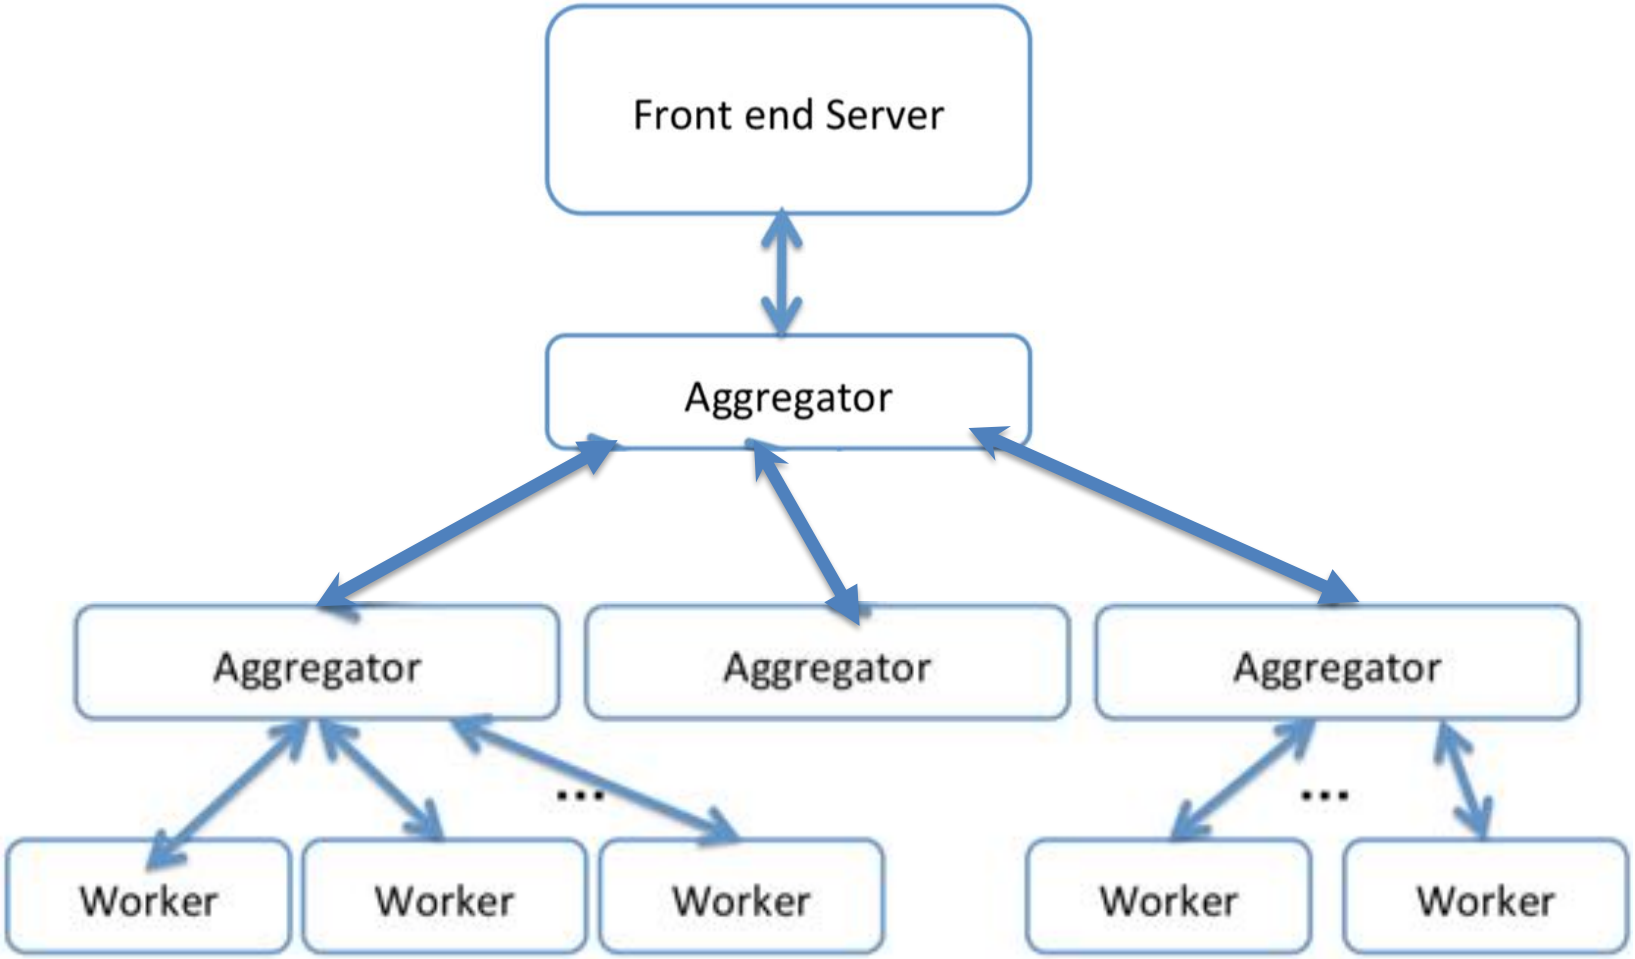

Data Center Network Design

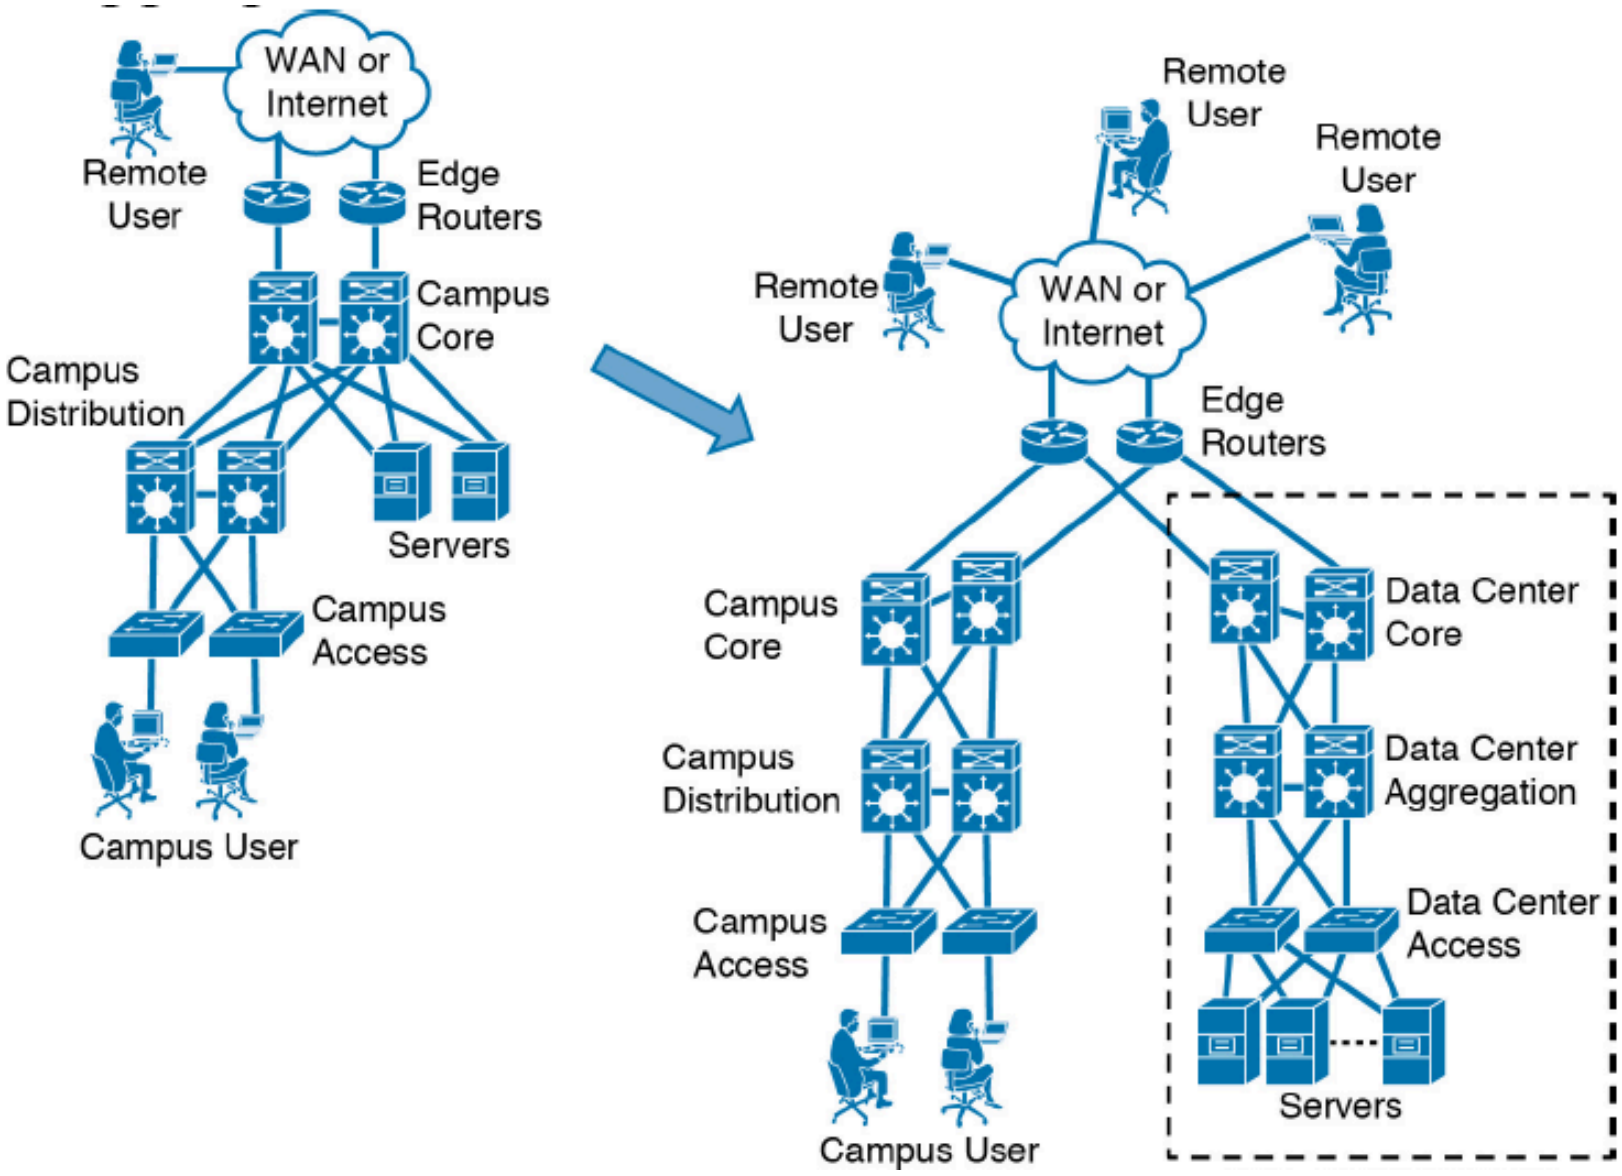

Data Center Network Design

- Three-tier Data Center Design: Access, Aggregation, Core
  - Each server connects to 2 access switches
  - Access switches connect to 2 aggregation switches
  - Aggregation switches connect to 2 core routers
  - Core routers connect to edge routers
  - Aggregation layer is the transition point between L2-switched access layer and L3-routed core layer

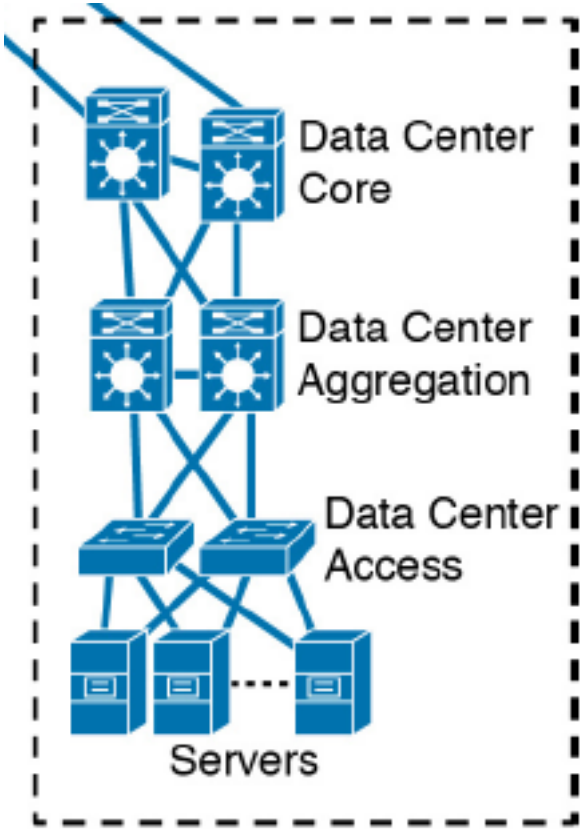

Data Center Network Design

- Example: **Facebook**
  - https://www.youtube.com/watch?v=mLEawo6OzFM

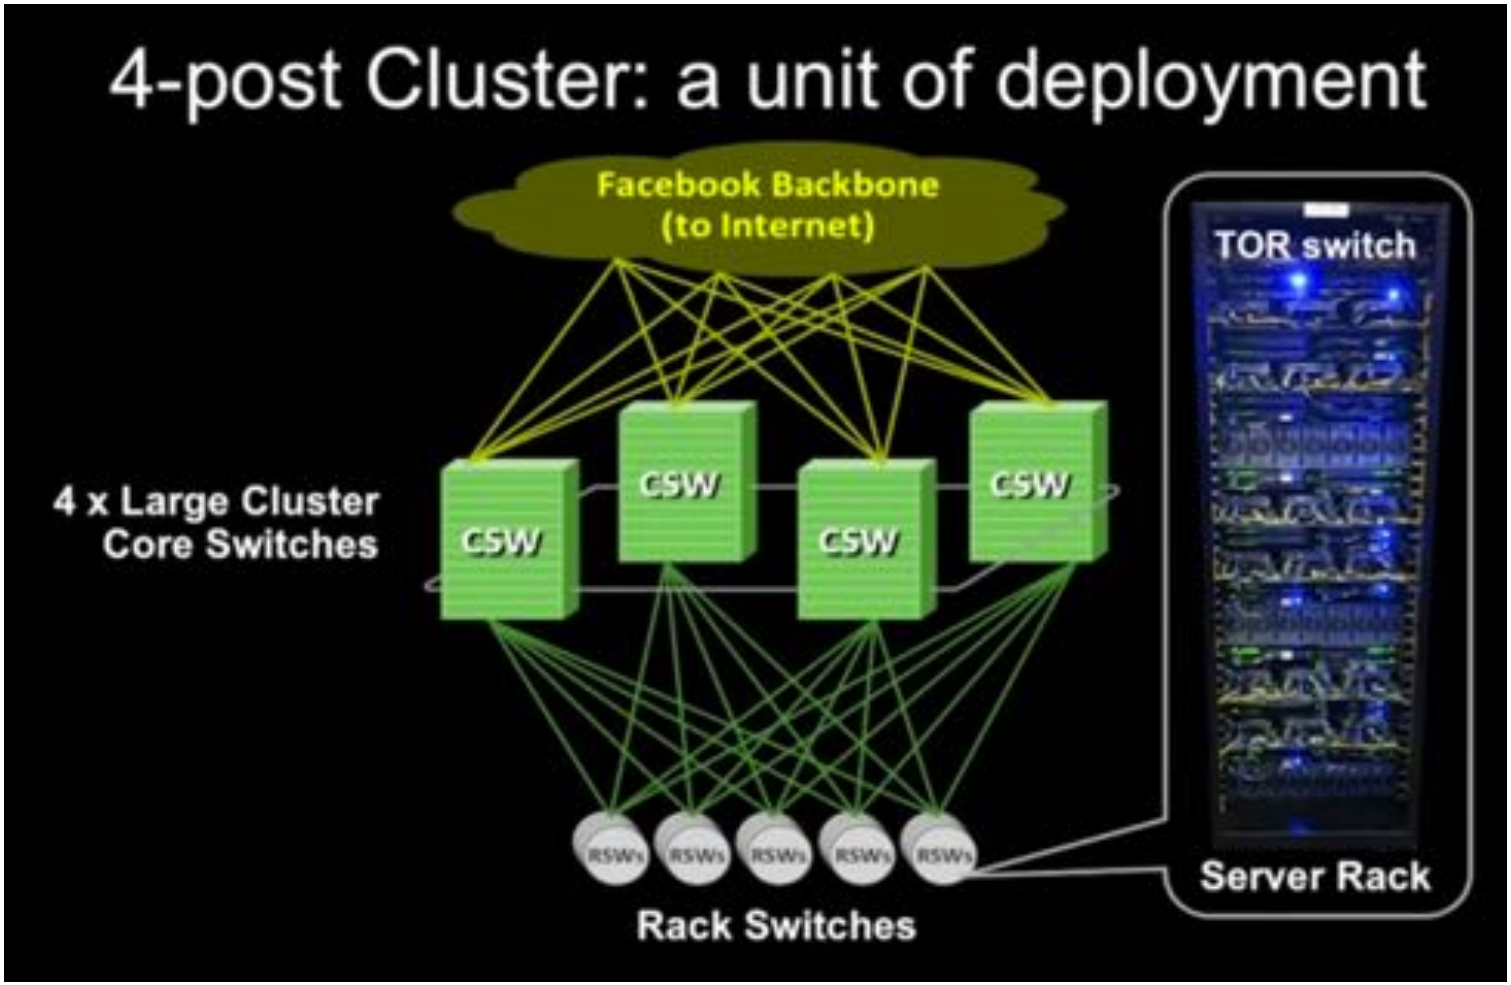

Data Center Network Design

- Example: **Facebook**
  - https://www.youtube.com/watch?v=mLEawo6OzFM

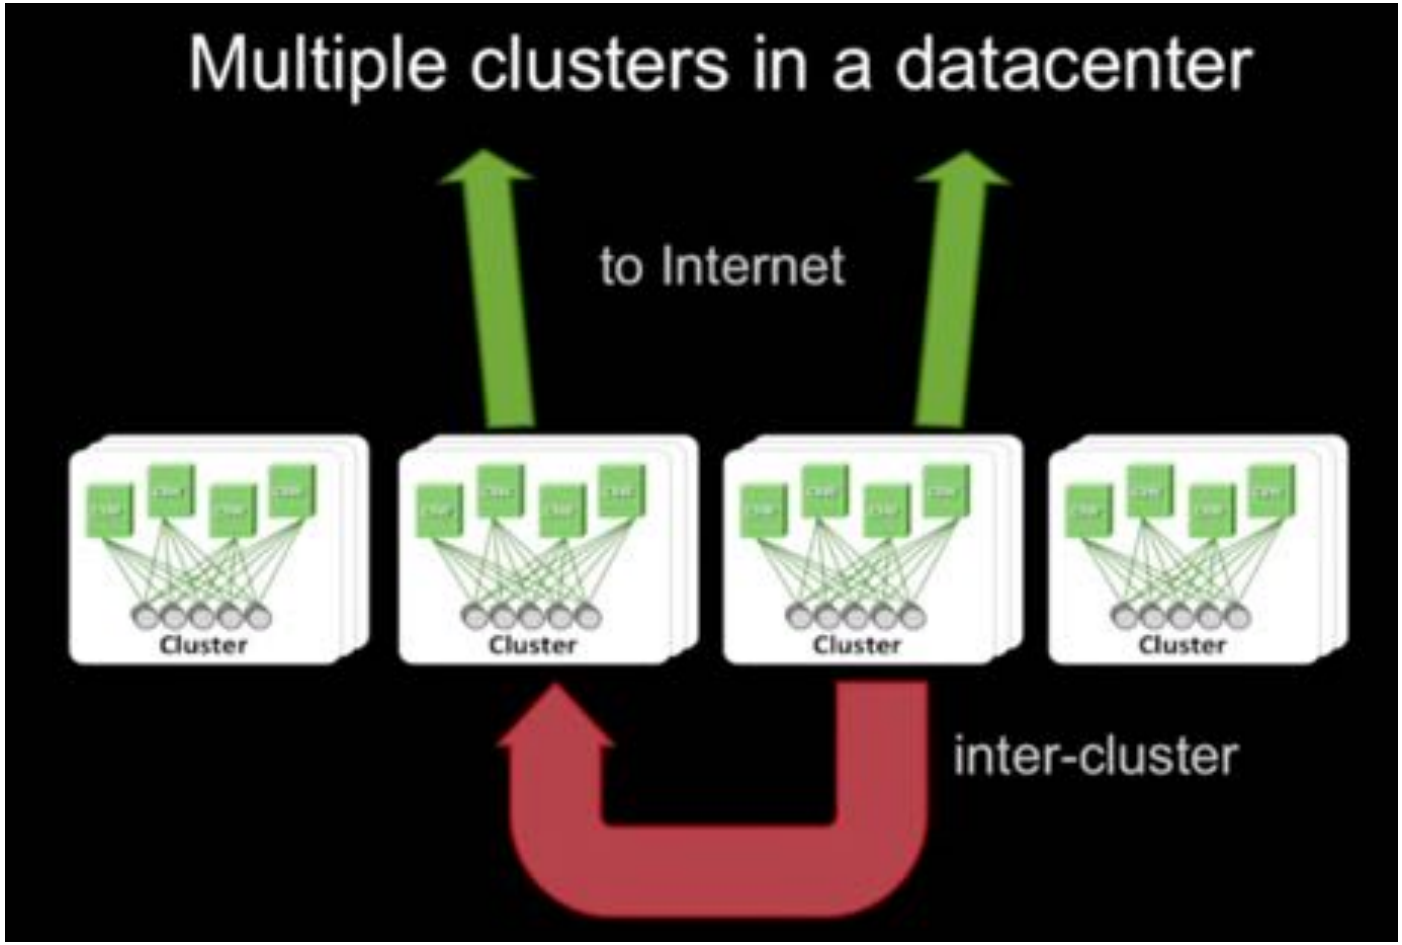

Data Center Network Design

- Example: **Facebook**
  - High inter-cluster traffic

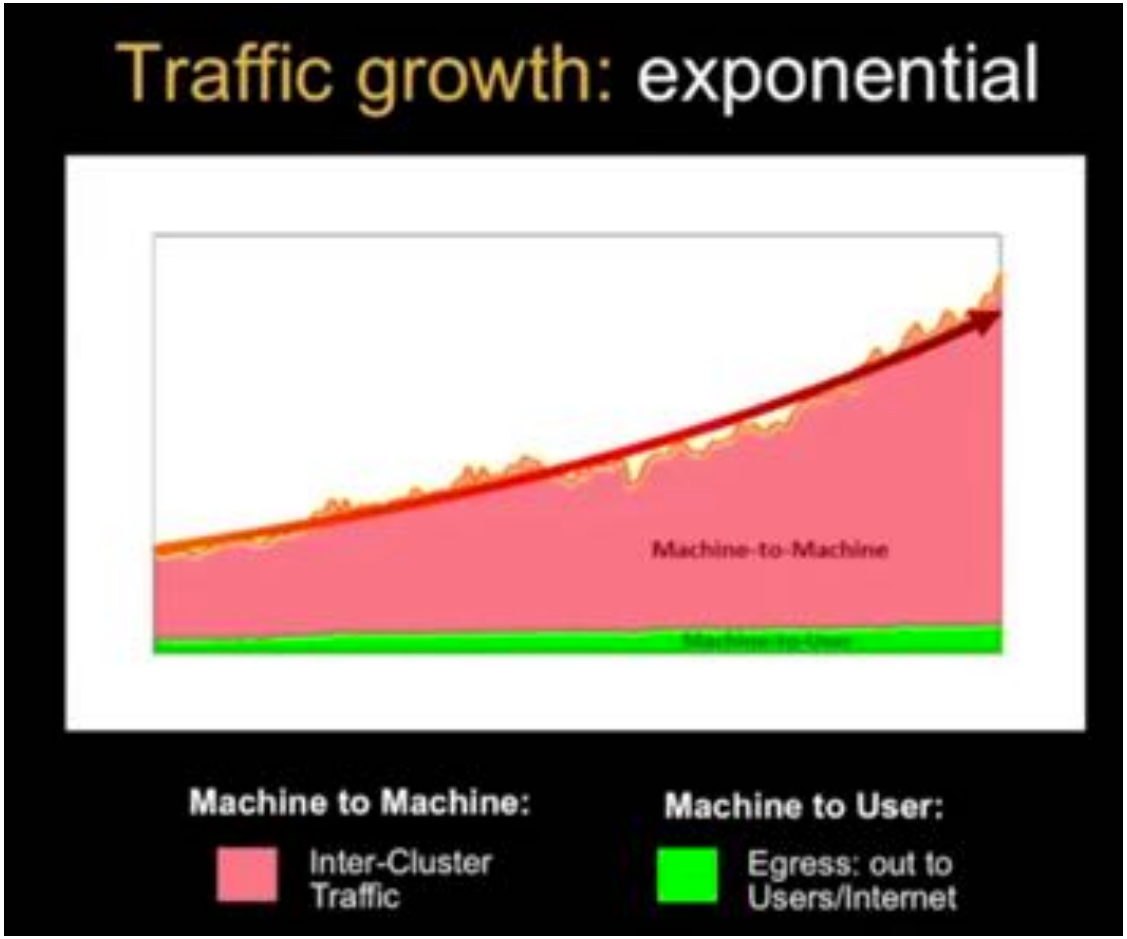

Data Center Network Design

- Example: **Facebook Fat-Tree Topology**

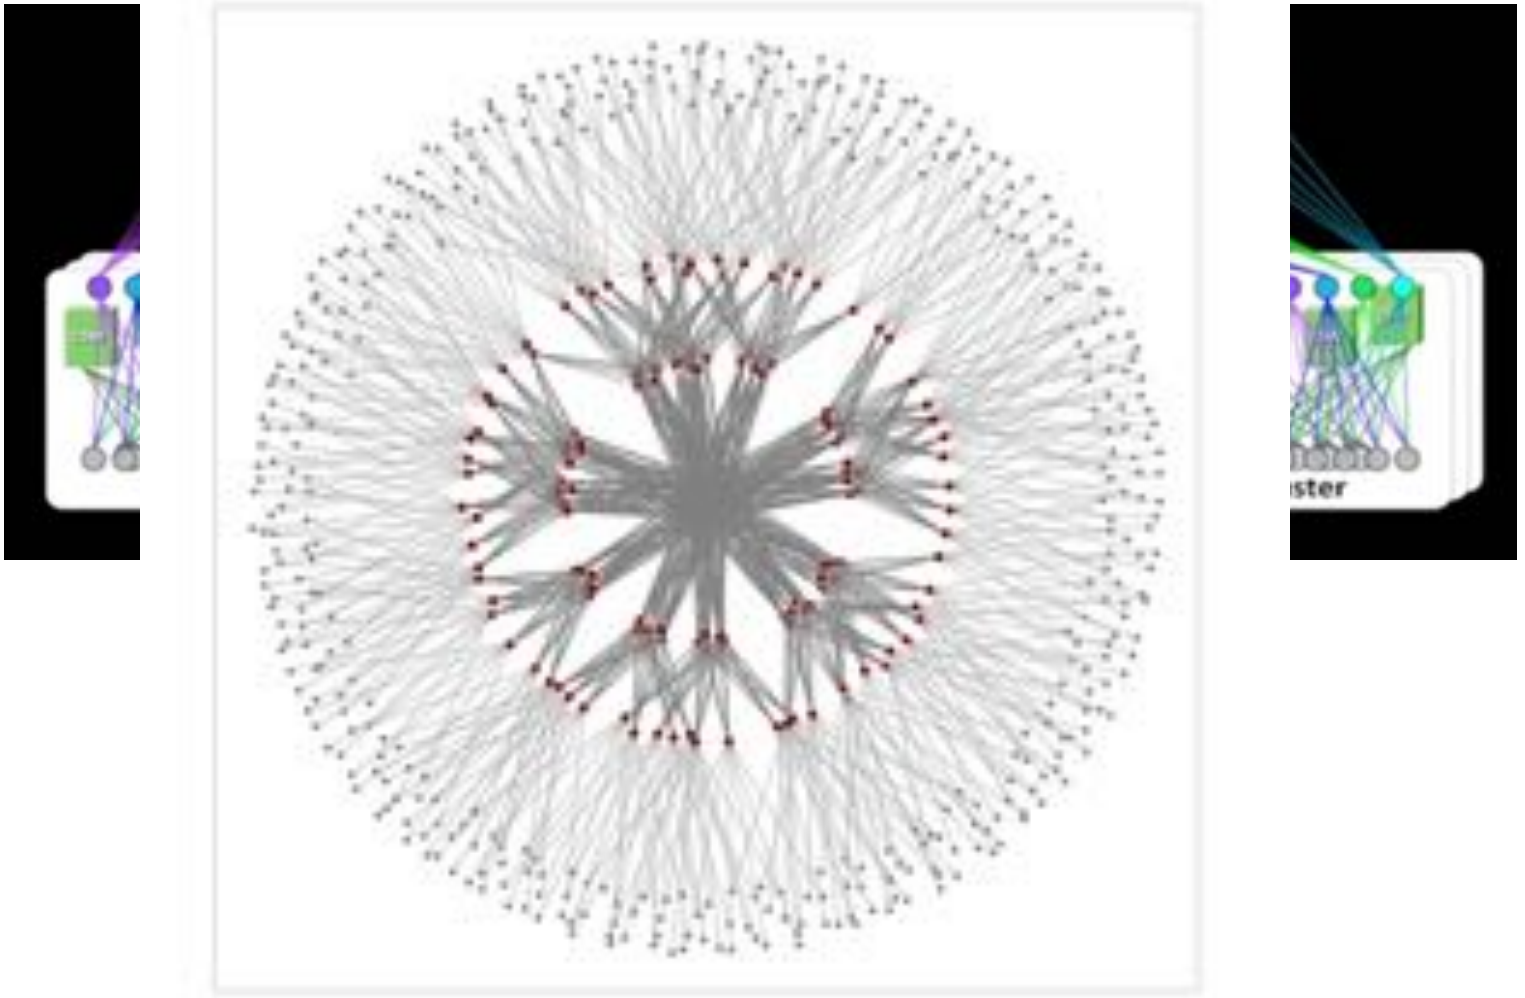

3-level fat-tree - 432 servers, 180 switches, degree 12

Data Center Network Design

- Example: **Facebook Fat-Tree Topology**
  - **Many paths**

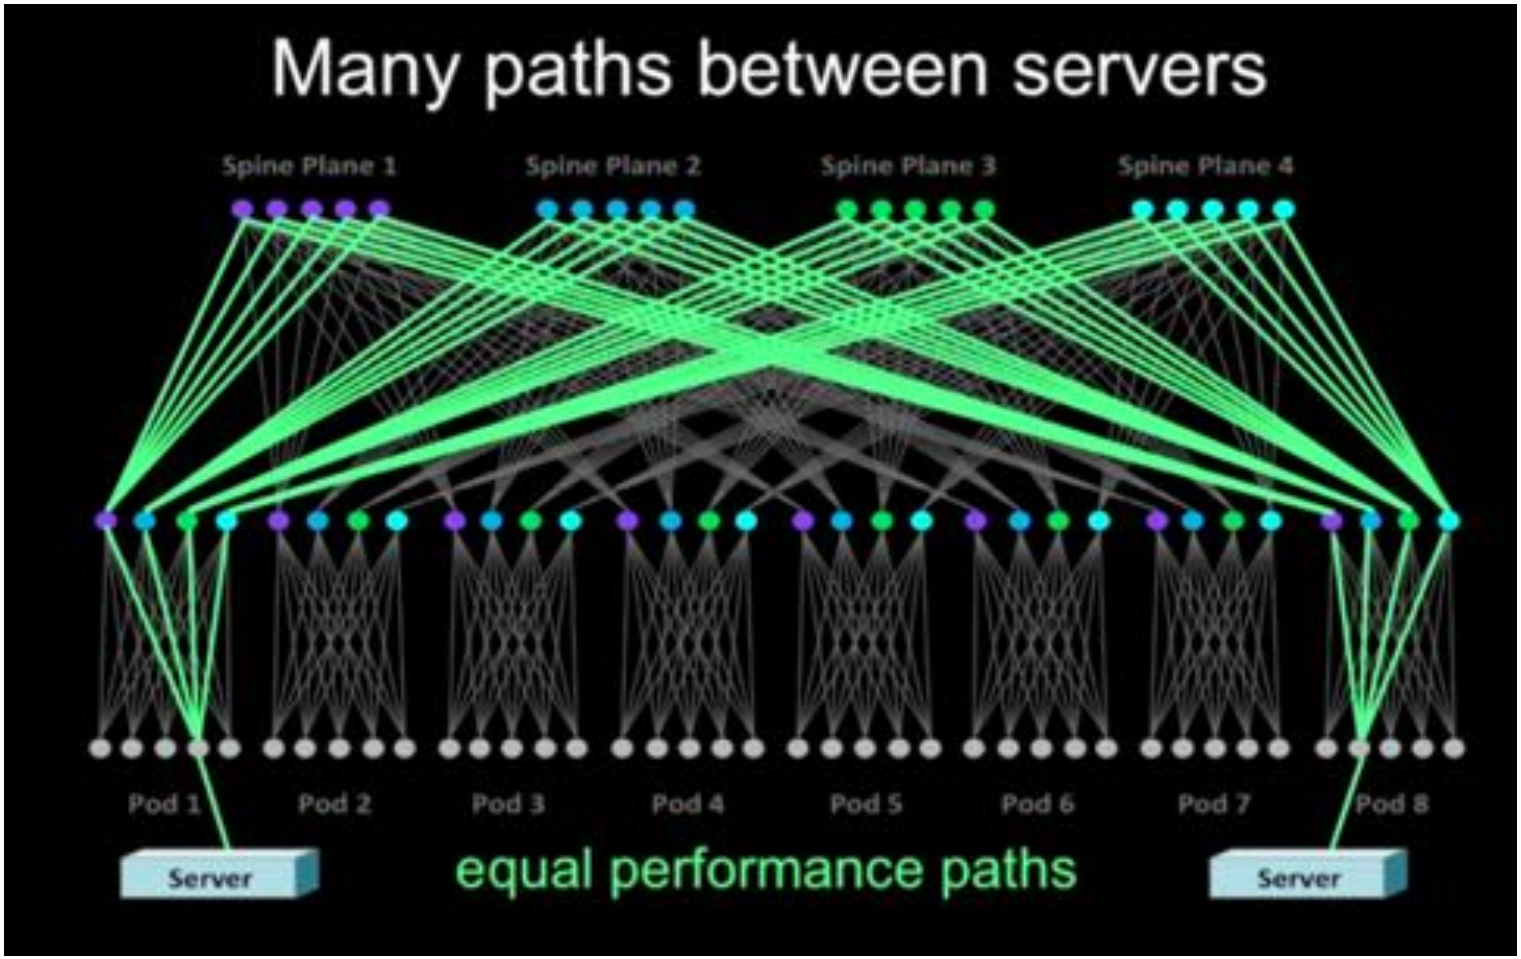

Data Center Network Design

- Example: **Facebook Fat-Tree Topology**
  - **Load balancing**

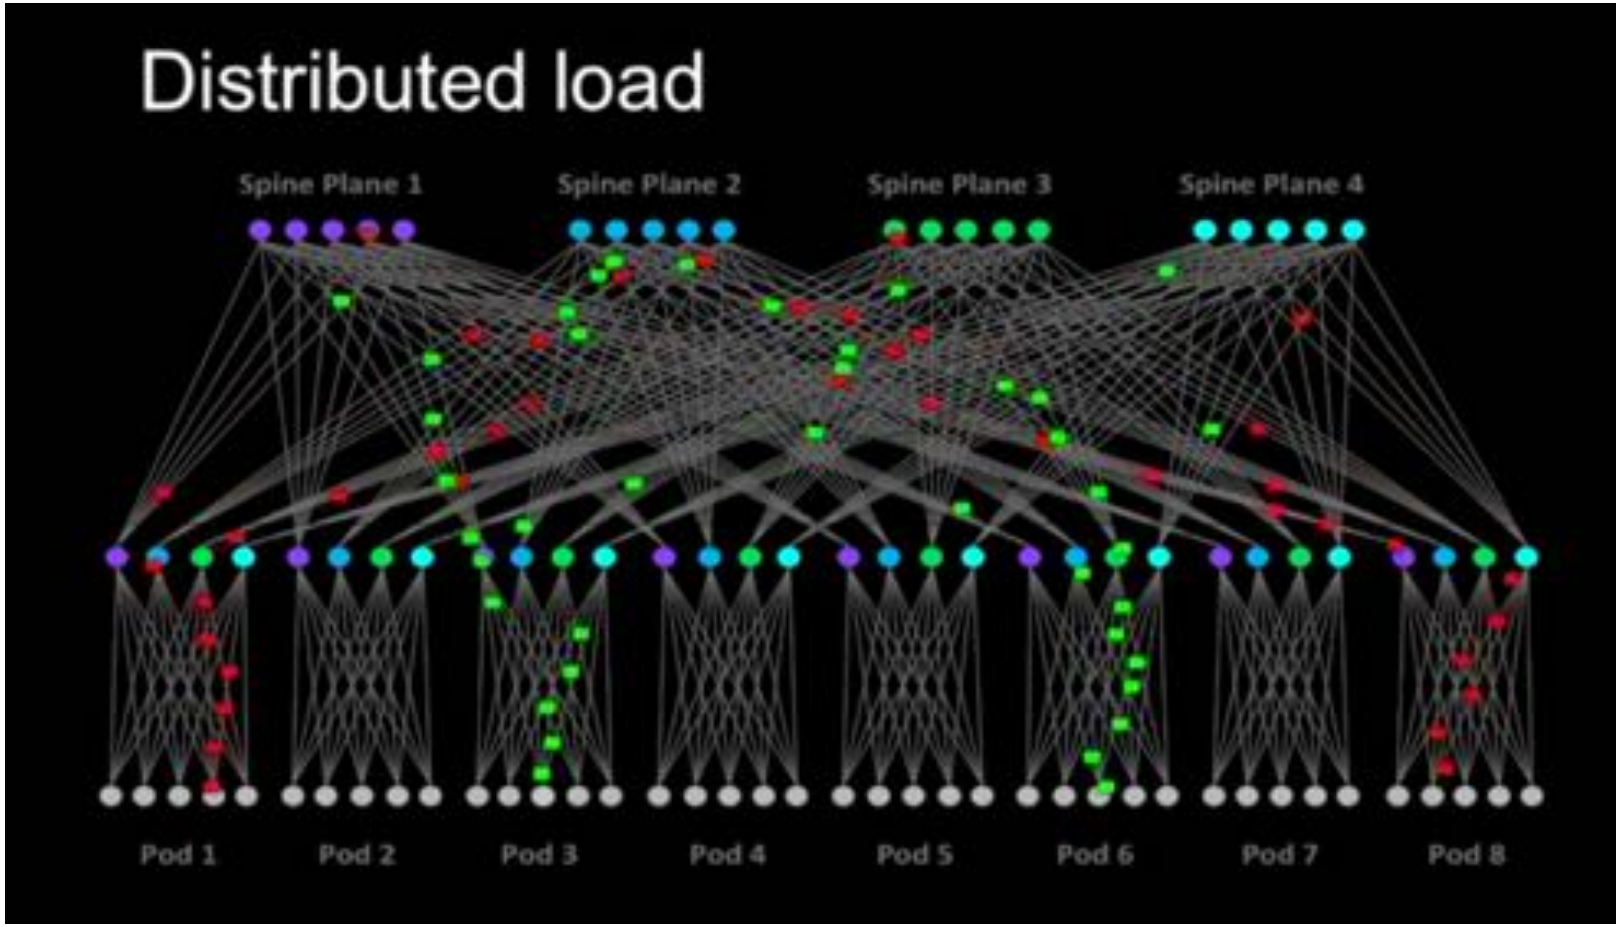

Data Center Network Design

- Example: **Facebook Fat-Tree Topology**
  - **Fault-tolerance**

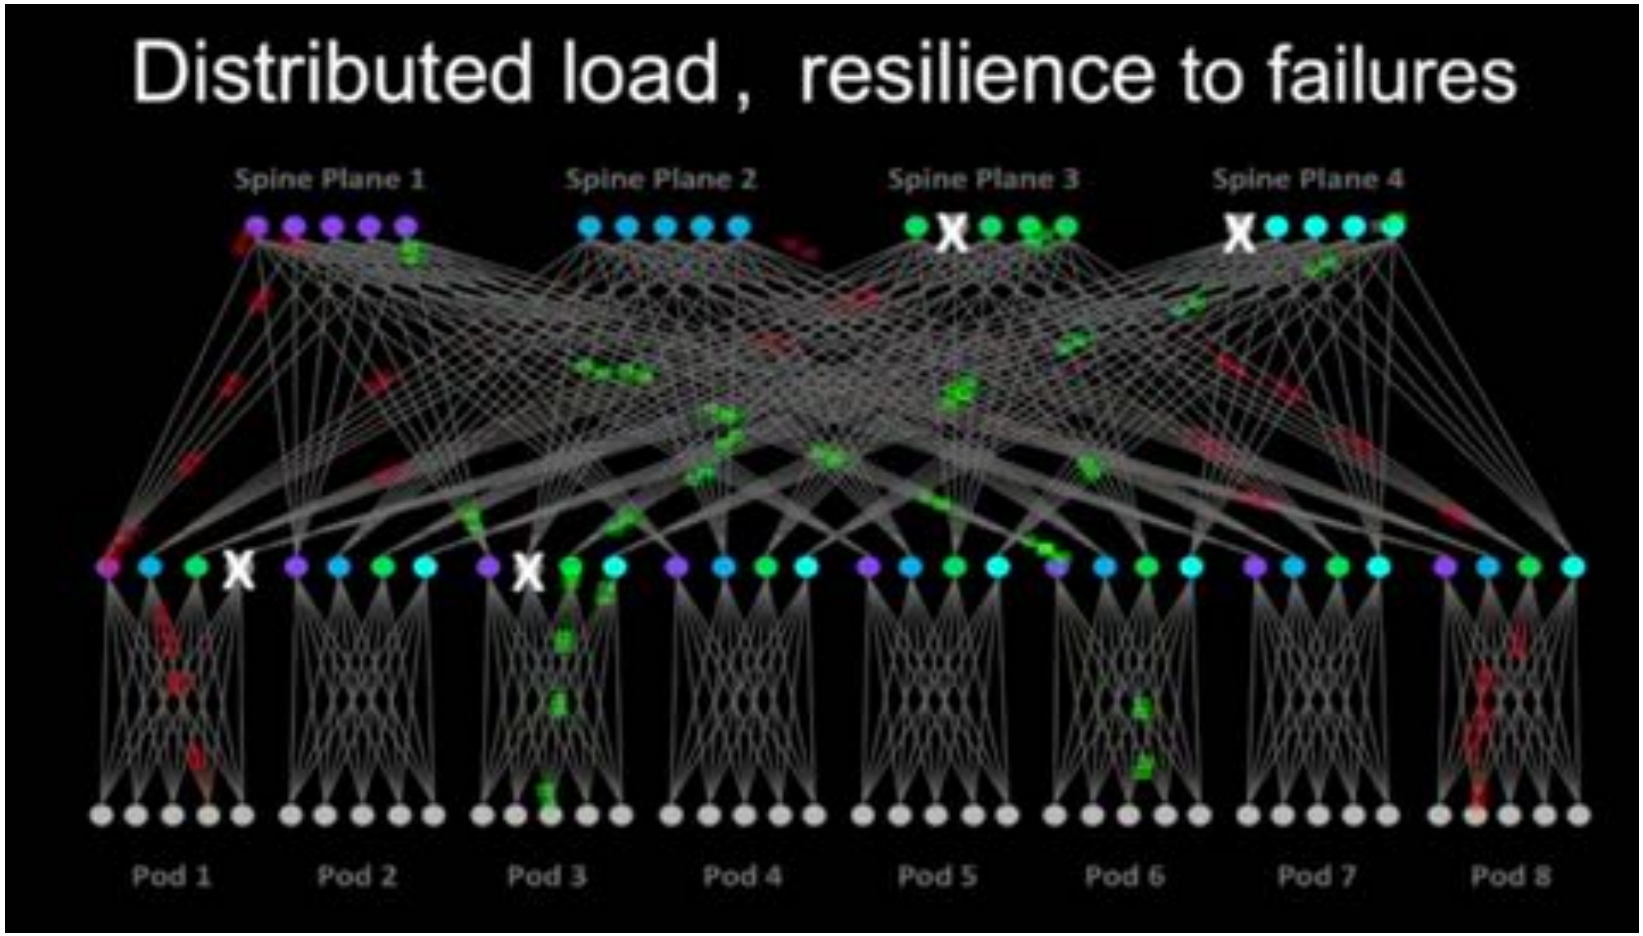

Data Center Network Design

- Example: **Facebook Fat-Tree Topology**
  - **Flexible, Scalable and Easy!**

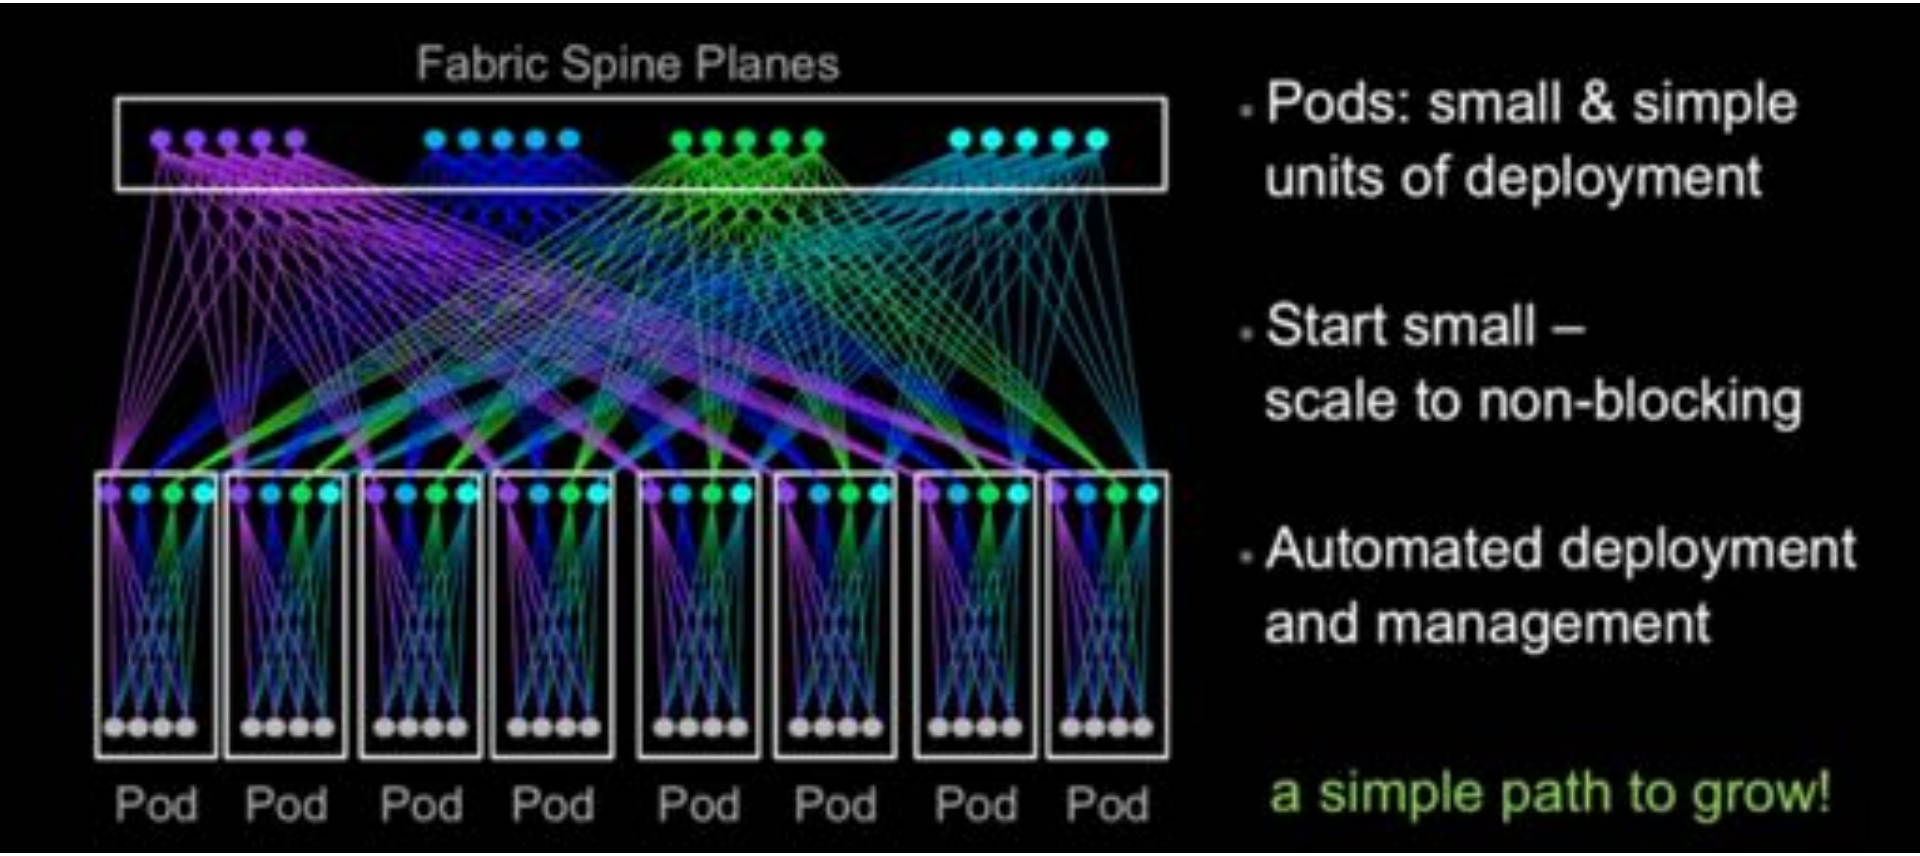

Data Center Network Design

- Fat-Tree (SIGCOMM '08): three-layer topology (edge, aggregation and core)
  - each pod consists of $(k/2)^2$ servers & 2 layers of k/2 k-port switches
  - each edge switch connects to k/2 servers & k/2 aggregation switches
  - each aggregation switch connects to k/2 edge & k/2 core switches
  - $(k/2)^2$ core switches: each connects to k pods

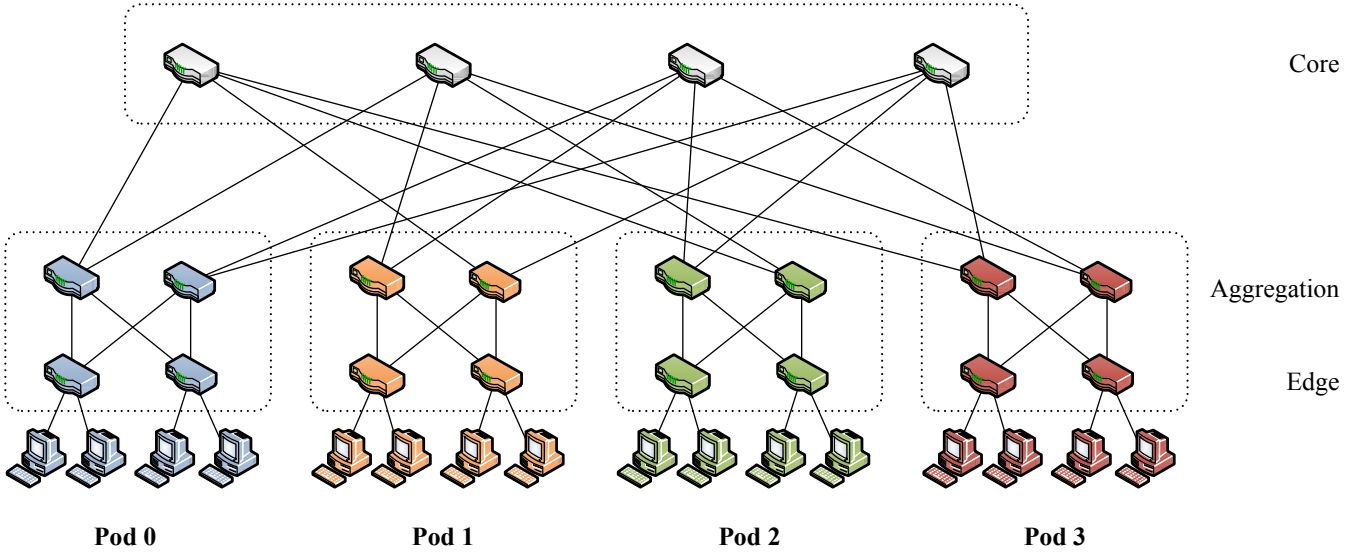

pods number: k=4

Data Center Network Design

- Recursive DCN Architecture (SIGCOMM '08):
  - A Level-0 subnet is the basic building block. It contains inter-connected servers. Each level-k subnet has multiple level-(k-1) subnets.

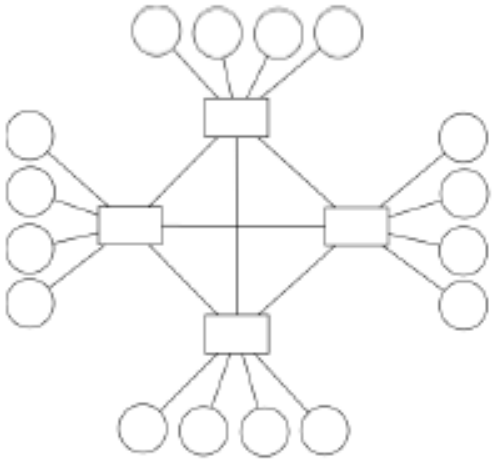

DCell-0

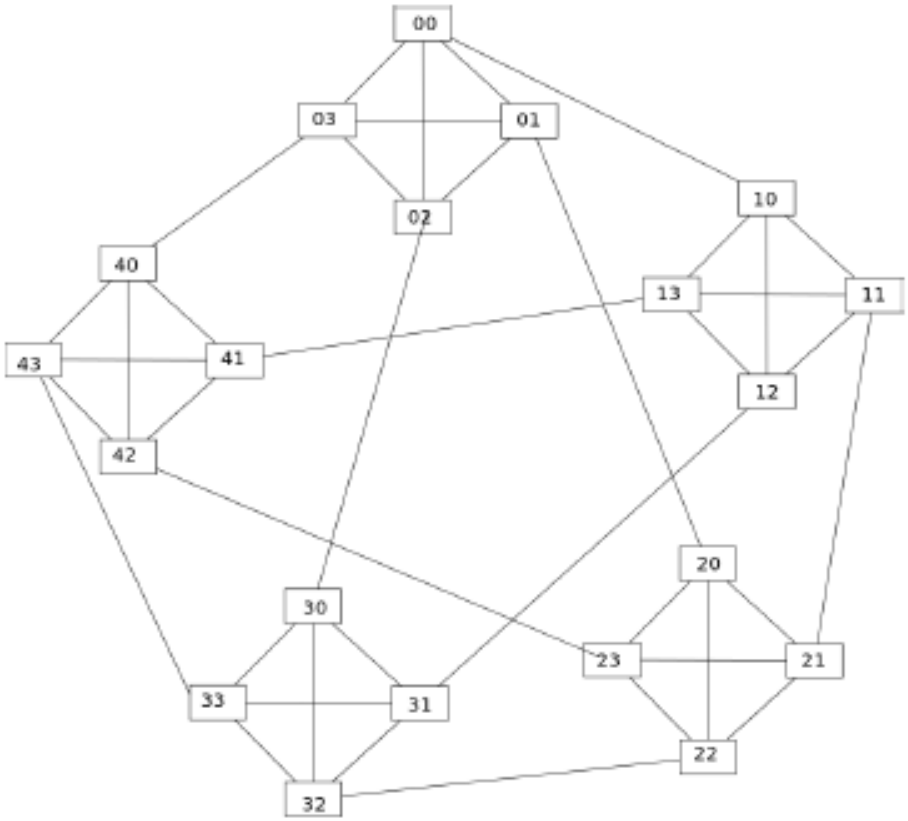

DCell-1 formed by 5 DCell-0

Data Center Network Design

- Bcube (SIGCOMM '09)
  - Servers with multiple ports
  - Switches that connect to a constant number of servers
  - A recursively defined structure: A $BCube_0$ is simply n server connecting to an n-port switch; A $BCube_1$ is constructed from n $BCube_0$s and n n-port switches.

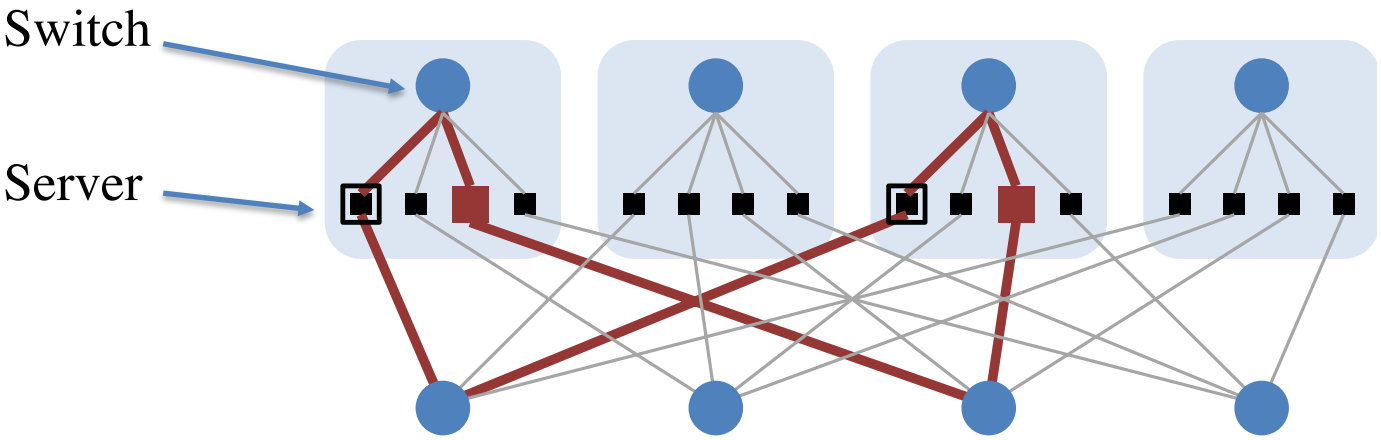

(b) BCube

Data Center Network Design

Energy: Results

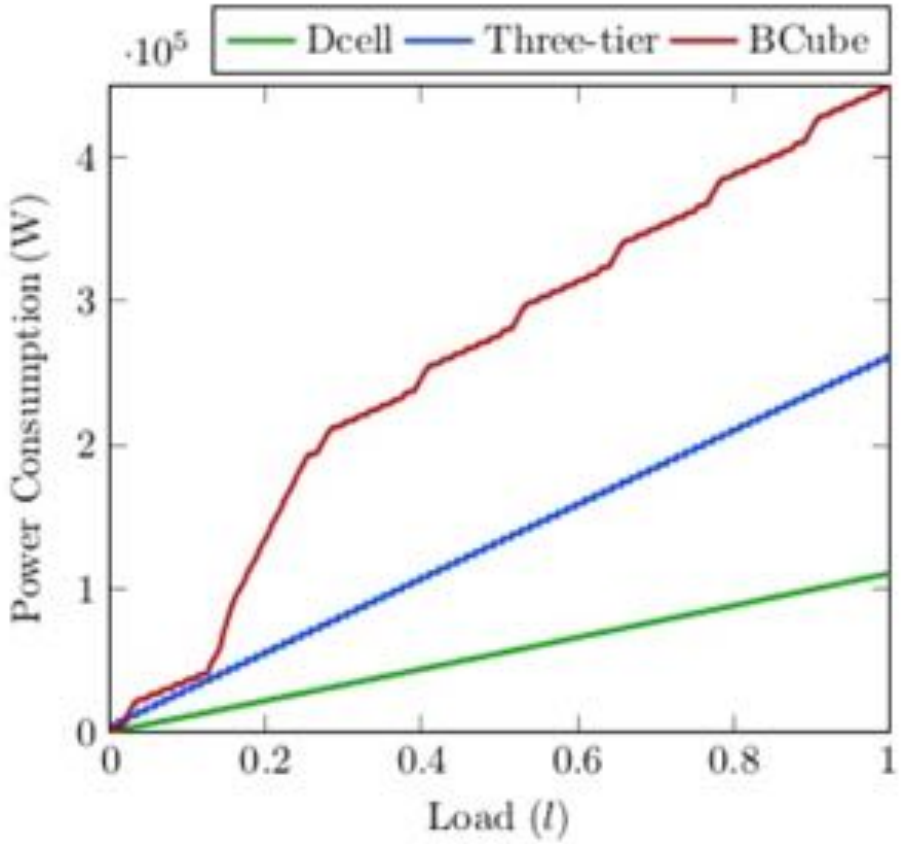

http://www.slideshare.net/cfiandra/dat-43495614

Data Center Network Design

Communication: Results

Performance Index	Architectures		
	Three-tier	BCube	DCell
Latency (40 B)	1.98 µs	3.93 µs	4.73 µs
Latency (1500 B)	28.34 µs	73.72 µs	93.92 µs
Hop distance	5.78	7.00	8.94
Server Degree Connectivity	1	4	2.79

http://www.slideshare.net/cfiandra/dat-43495614

References

- http://www.cse.wustl.edu/~jain/cse570-13/ftp/m_02trn.pdf
- http://www.cse.wustl.edu/~jain/cse570-13/ftp/m_03dct.pdf
- Data Center Networking – Connectivity and Topology Design Guide, enterasys secure networks, http://www.enterasys.com/company/literature/datacenter-design-guide-wp.pdf