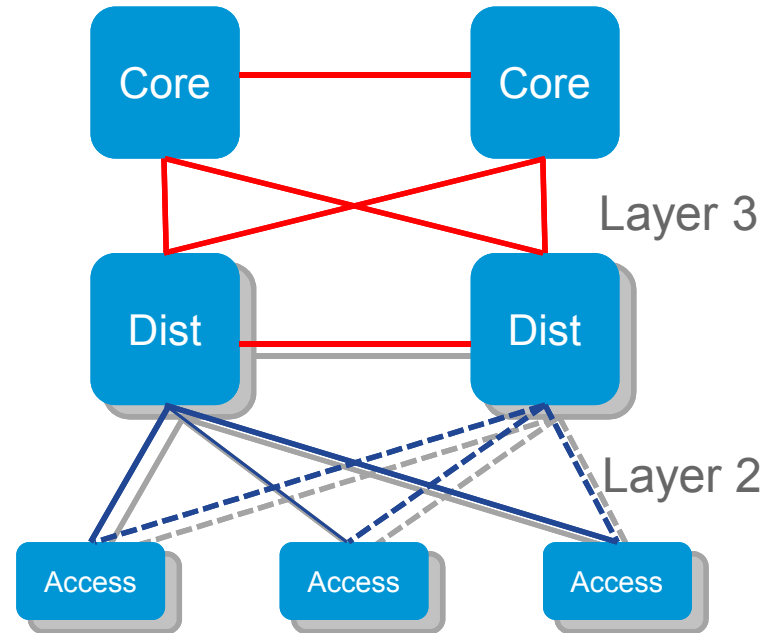


DC Architecture

- Spine/Leaf vs. traditional
- How to get to spine/leaf
- Cooling consideration
- Racking/Stacking/Naming
- System Management
- Optics and Cabling
- Power Considerations


Traditional Networking Model

- Addresses needs based on constraint
- Cost is an important factor
- Oversubscription is built into the model
- Built for North/South traffic patterns
- Mature, well understood
- Other design advantages as well



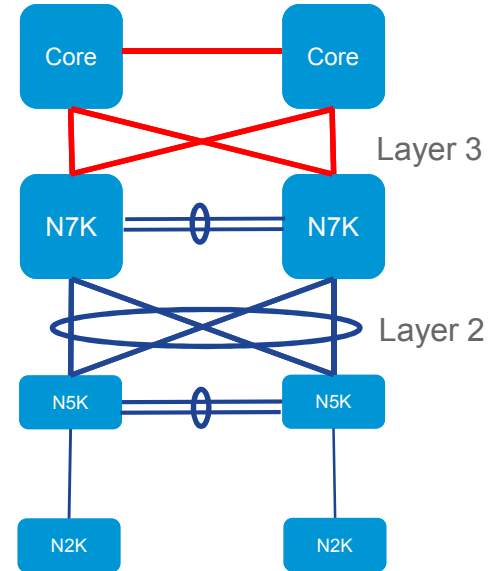
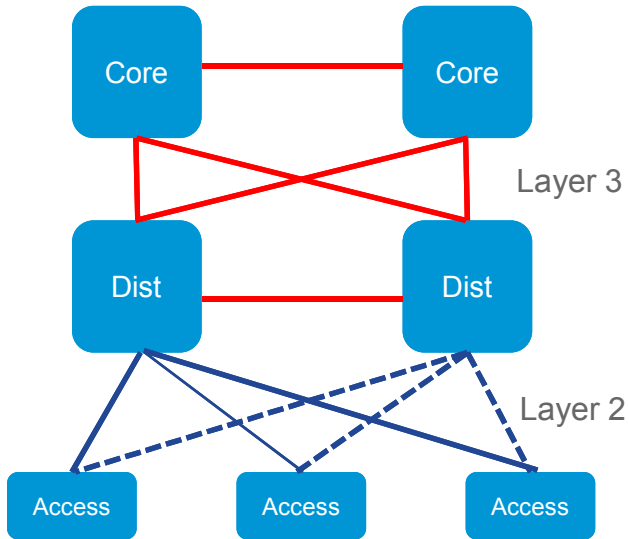
Transitioning into the DC

- In the DC...

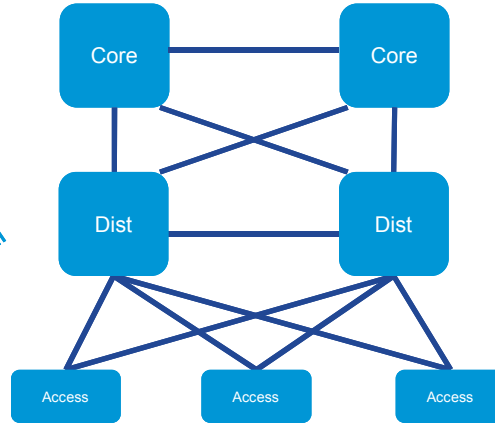


- We've taken the same data flow model and overlaid it on the DC
- Some differences:
 - Less Over-subscription
 - Need to look at East-West traffic (in addition to North-South)
 - “Higher” utilization and bursty traffic may require deeper buffers

How We Designed DC Switching 10 Years Ago



Evolution of Network Requirements



Campus Reqs

- BYOD/Mobility
- Voice and Video
- PoE→PoE+→UPOE...
- QoS
- VDI
- L2 → L3 Access
- 1G → 2.5G Access Ports
- 1G → 10G → 40G uplink
- SDN (OF/ONE PK/APIC-EM)
- IOT


Data Center Reqs

- Distributed workloads
- Virtualization (VM → Cont → MS)
- L2 Explosion
- 10/25G to hosts
- 40/100G uplinks
- East West Traffic (Layer 2)
- Multi DC
- Automation
- Hybrid Cloud
- SDN

...Evolution of Network Architectures

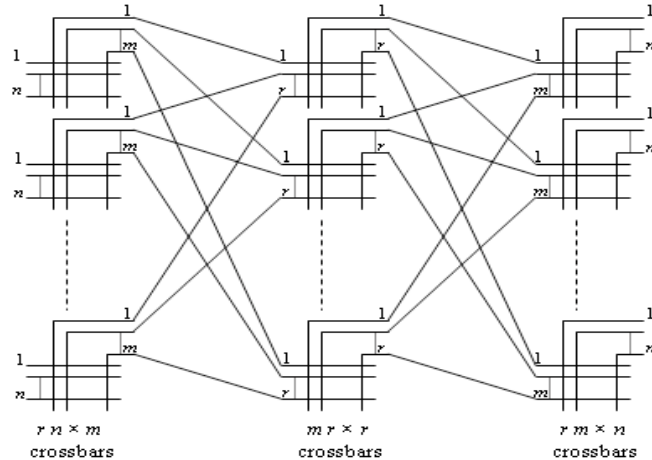
The Network Fabric

- Architecturally, a Spine and Leaf topology
- Internally, a protocol ties the elements together (underlay)
- Externally




The Roots (or tree) behind Network Fabrics

- Clos Network (Fabric)




- Fat Tree




Clos Network

- 1952-3: Charles Clos, researcher at Bell Labs published a paper describing non-blocking and survivable method to switch telephone calls
- Consists of multiple junctions/crossbars where traffic can flow
- Consists of 3 stages: Ingress, middle and egress
- Each junction in the stage is a crossbar switch
- Allows for full bandwidth between any 2 points – with resiliency




Fat Tree

- 1985: Charles E. Leiserson of MIT
- Tree Topology (Core/Dist/Access), which has "*skinny*" links
 - **Oversubscribed**
- The links in a fat-tree become "*fatter*" as one moves up the tree towards the root
 - **Non-oversubscribed**



Evolution of Network Architecture

Layer 3 Spine/Leaf Underlay Network

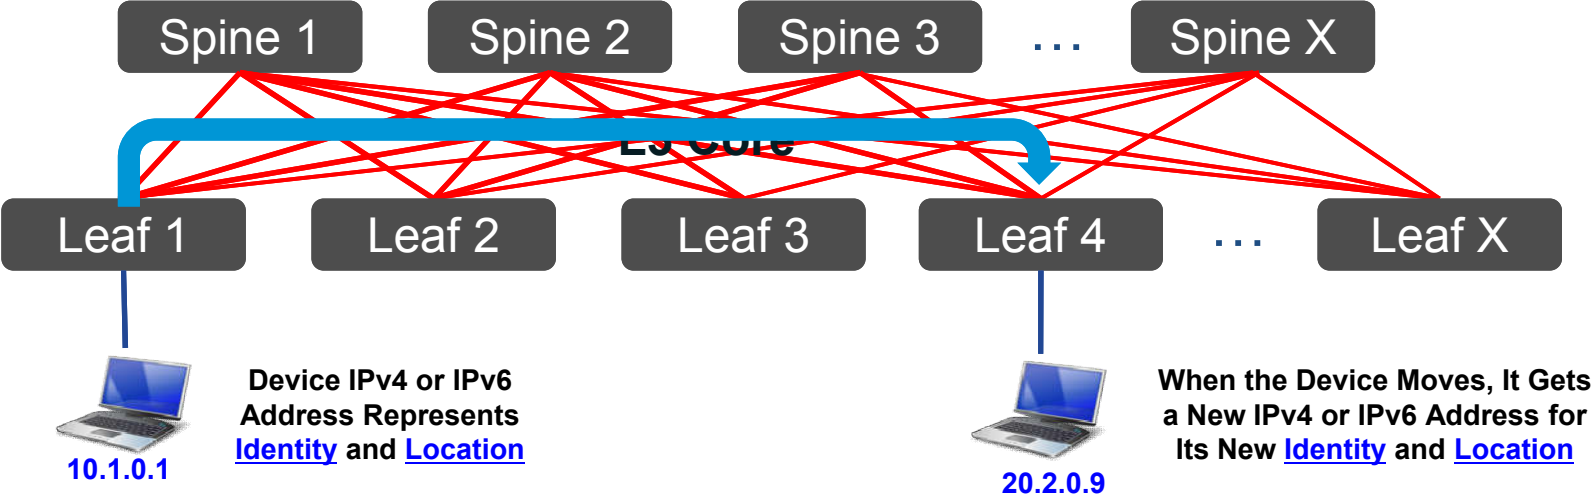


Flat, Fast and Flexible

- Any-to-any
- High Bandwidth
- Low Latency
- High Resiliency

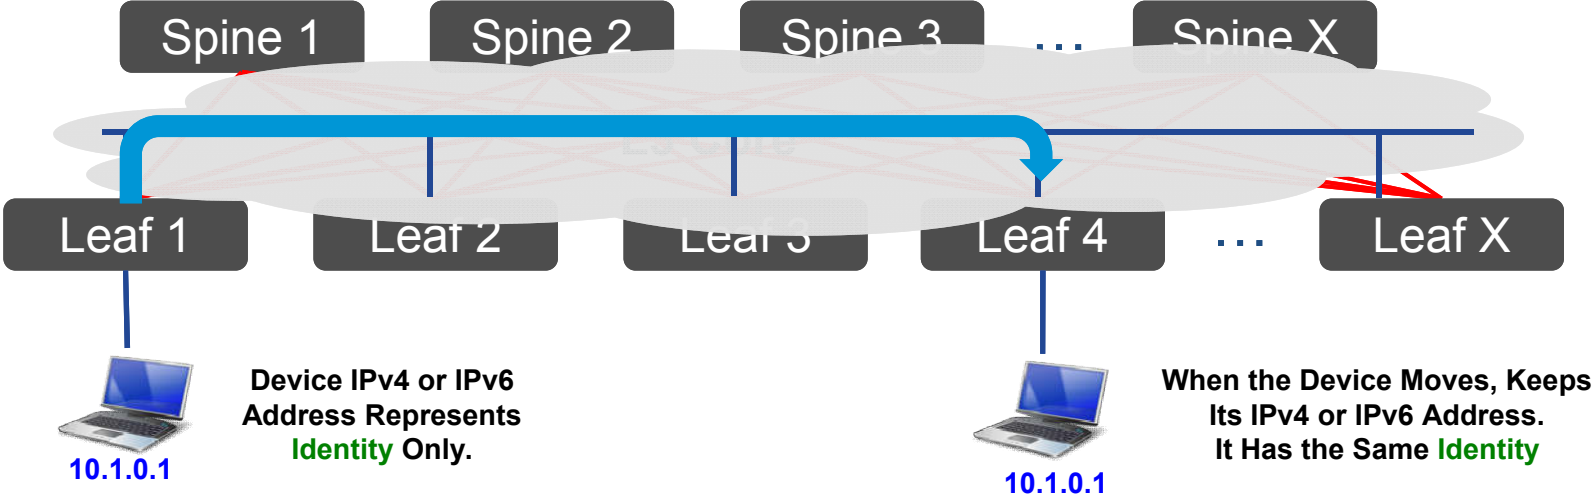
Evolution of Network Architecture

Layer 3 Spine/Leaf Underlay Network

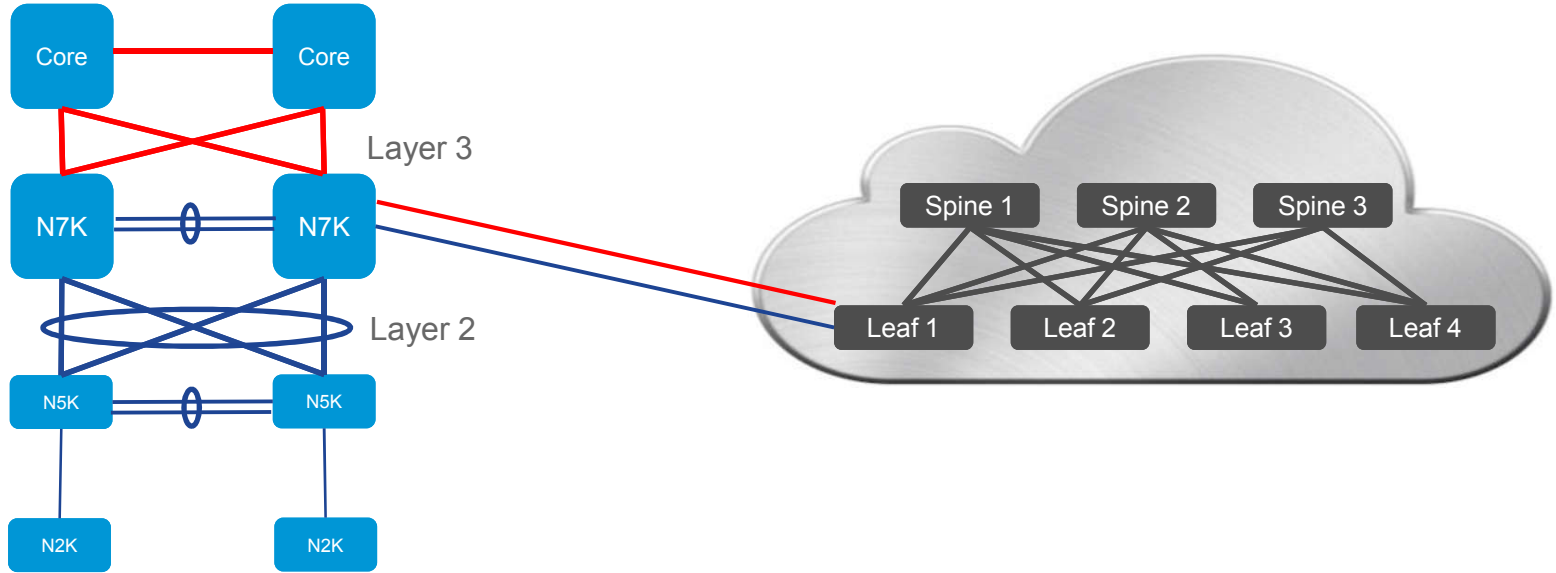


Evolution of Network Architecture

Layer 2/3 Overlay Network

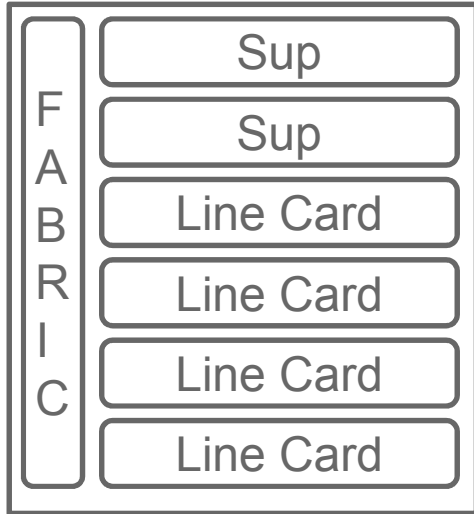


Transitioning to Spine/Leaf

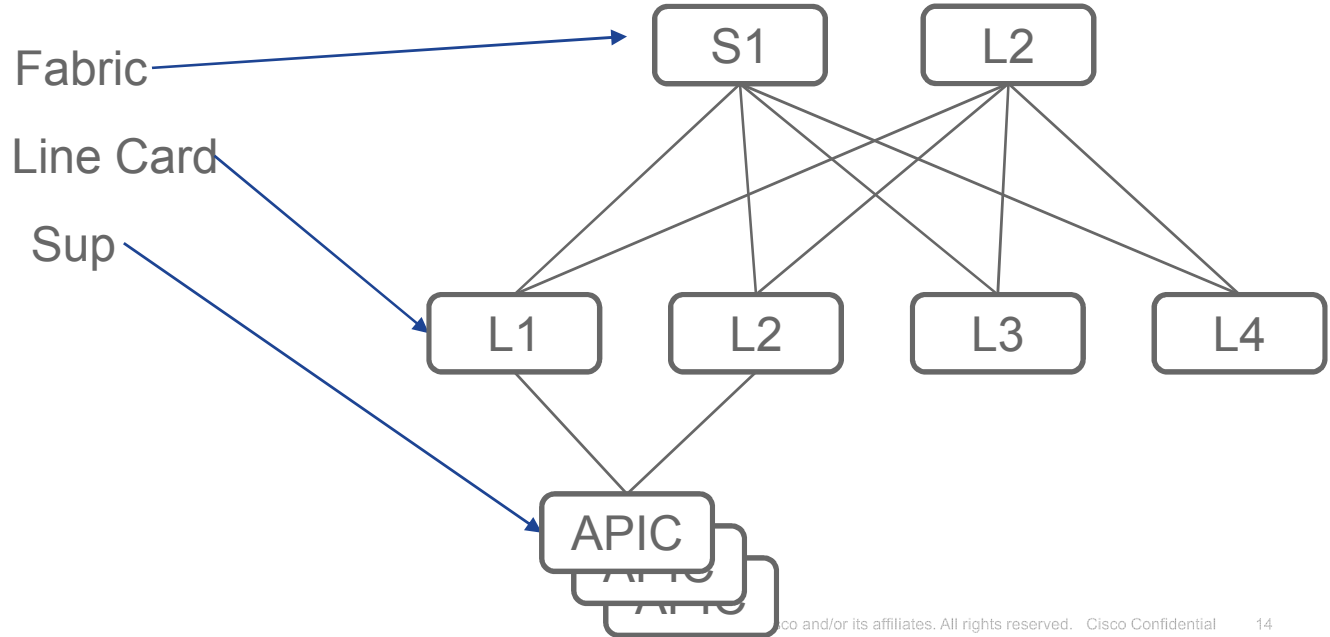


Chassis to Fabric Comparison

Traditional Catalyst



ACI Fabric



Racking/Stacking/Numbering

- Port side intake or exhaust?
- Gaps between equipment?
- Numbering Conventions?
- Safety?

DC Cooling Considerations

- Colo or Private DC?
- HVAC?
- Airflow?
- Equipment length?
- Blanking panels

System Management

- IP OOB/Dedicated Management Network
- Console Access
- Modem?

Optics and Cabling

- Know what type of fiber you have
- Know what type of fiber your ISP provides
- Are you in a CoLo? What is their protocol for MeetMe?
- What kind of structured cabling do you need?
- Easier to pull once, than send someone back out.
- Remember to come up with a standard on labeling

Power Considerations

- Always have a minimum of two power circuits
- Redundant power versus N+1
- If you have PDU's, consider 2M jumpers for power cables
- Data sheets are your friend!