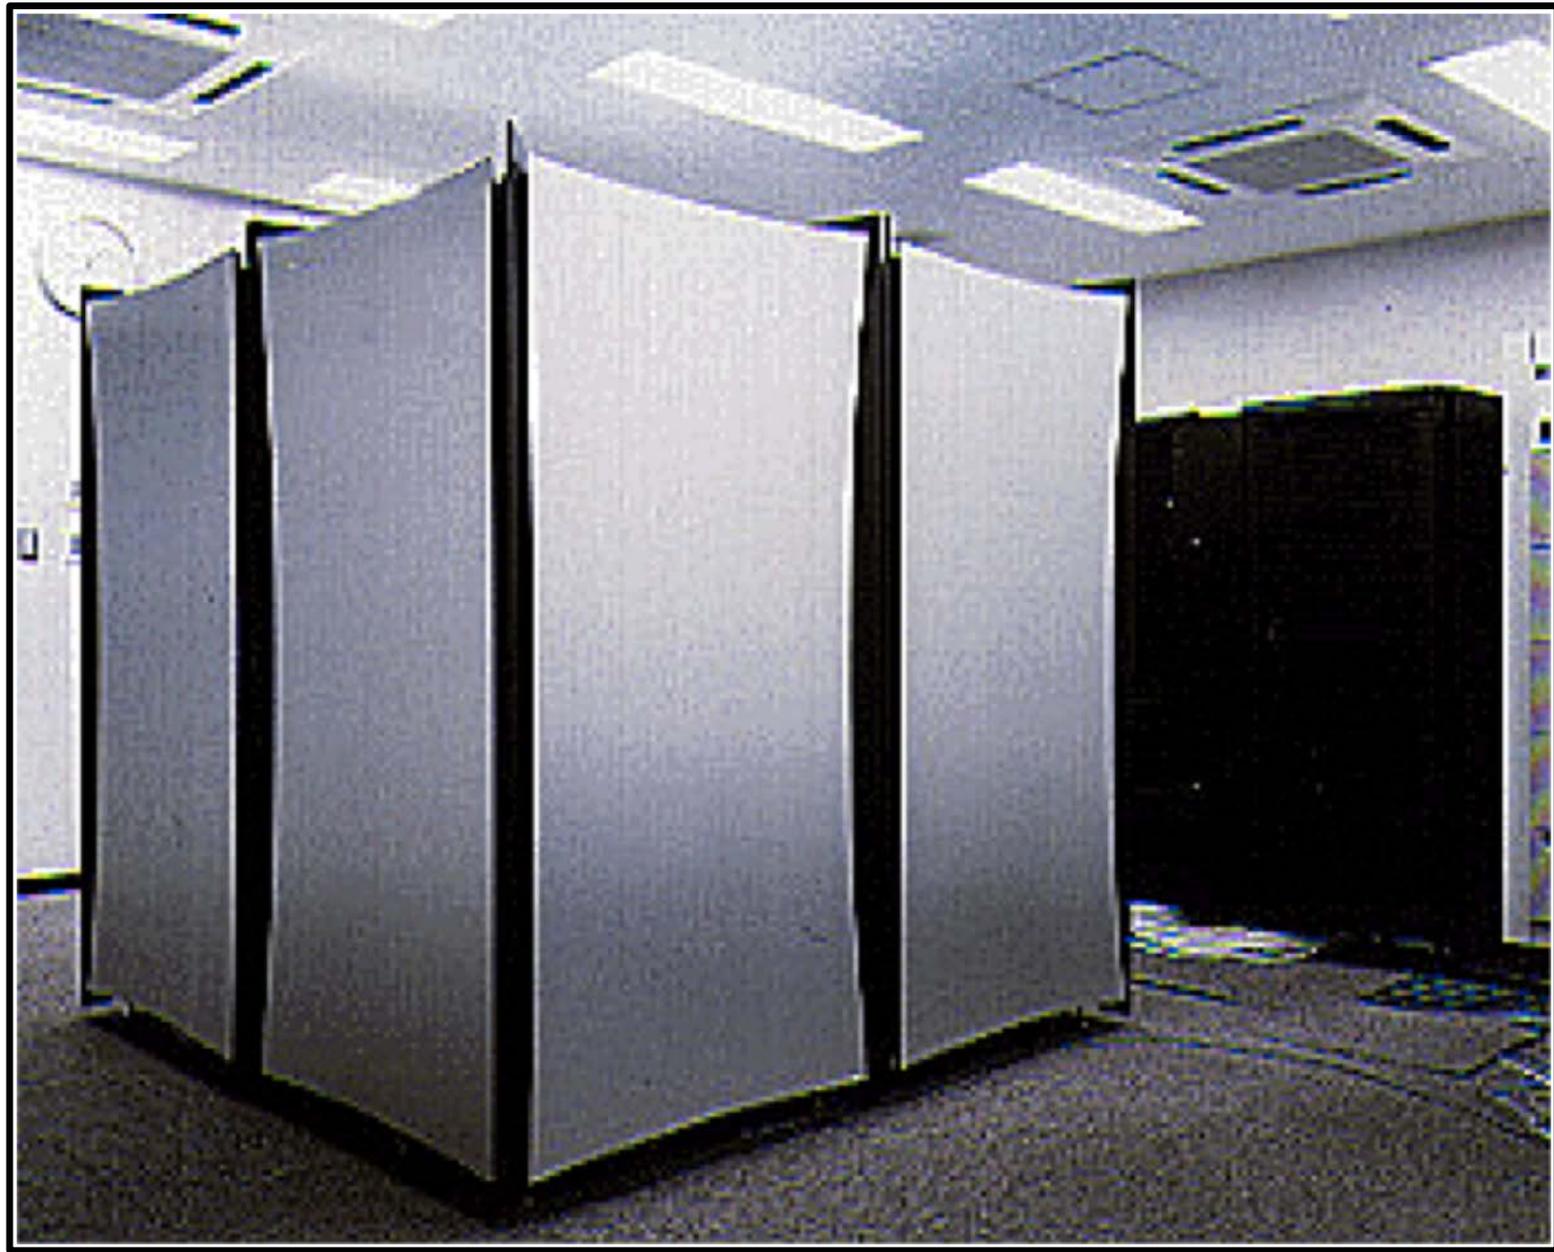




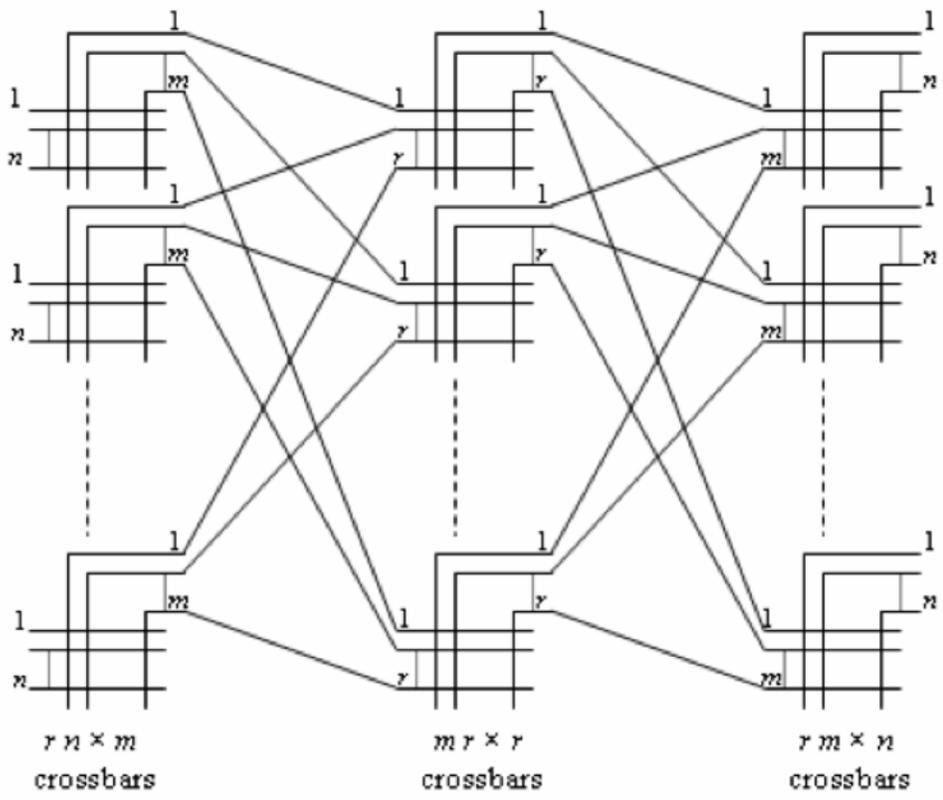
facebook
INFRASTRUCTURE

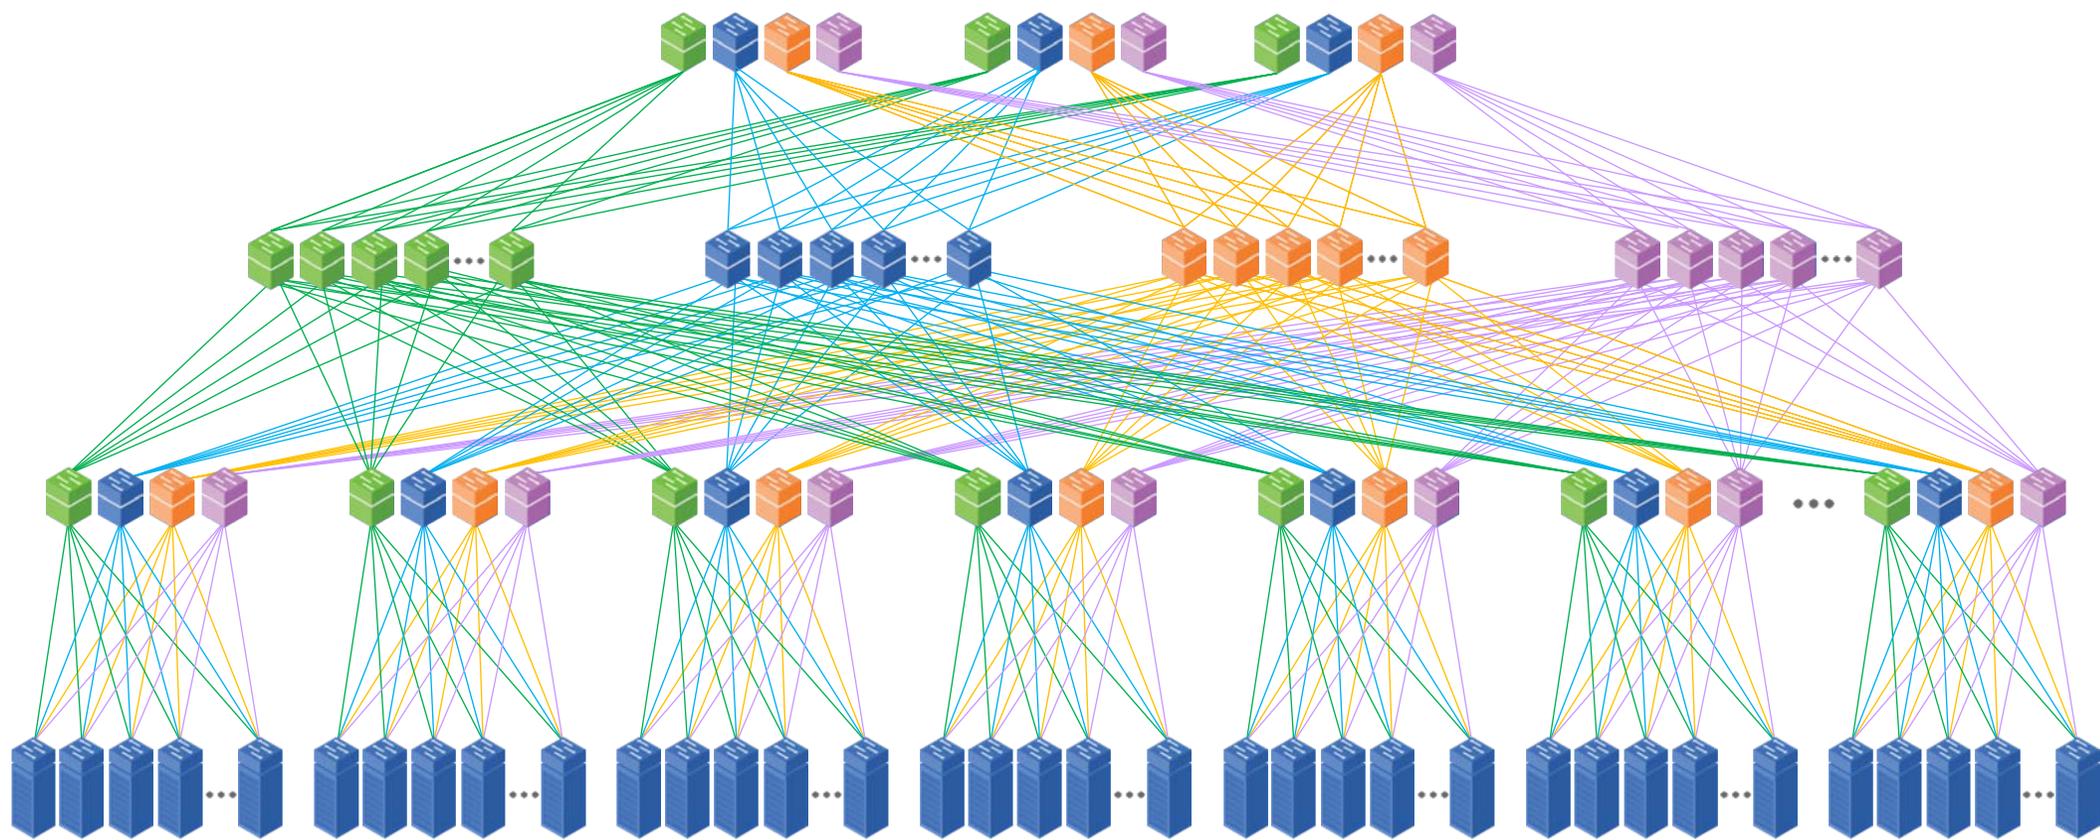




CLOS Topologies

Harnessing the Bleeding Edge of 1950s Telephone Switch Technology







Wedge 40

2014



Wedge 100

2015



Wedge 100S

2017

2015

6-pack



2016

Backpack

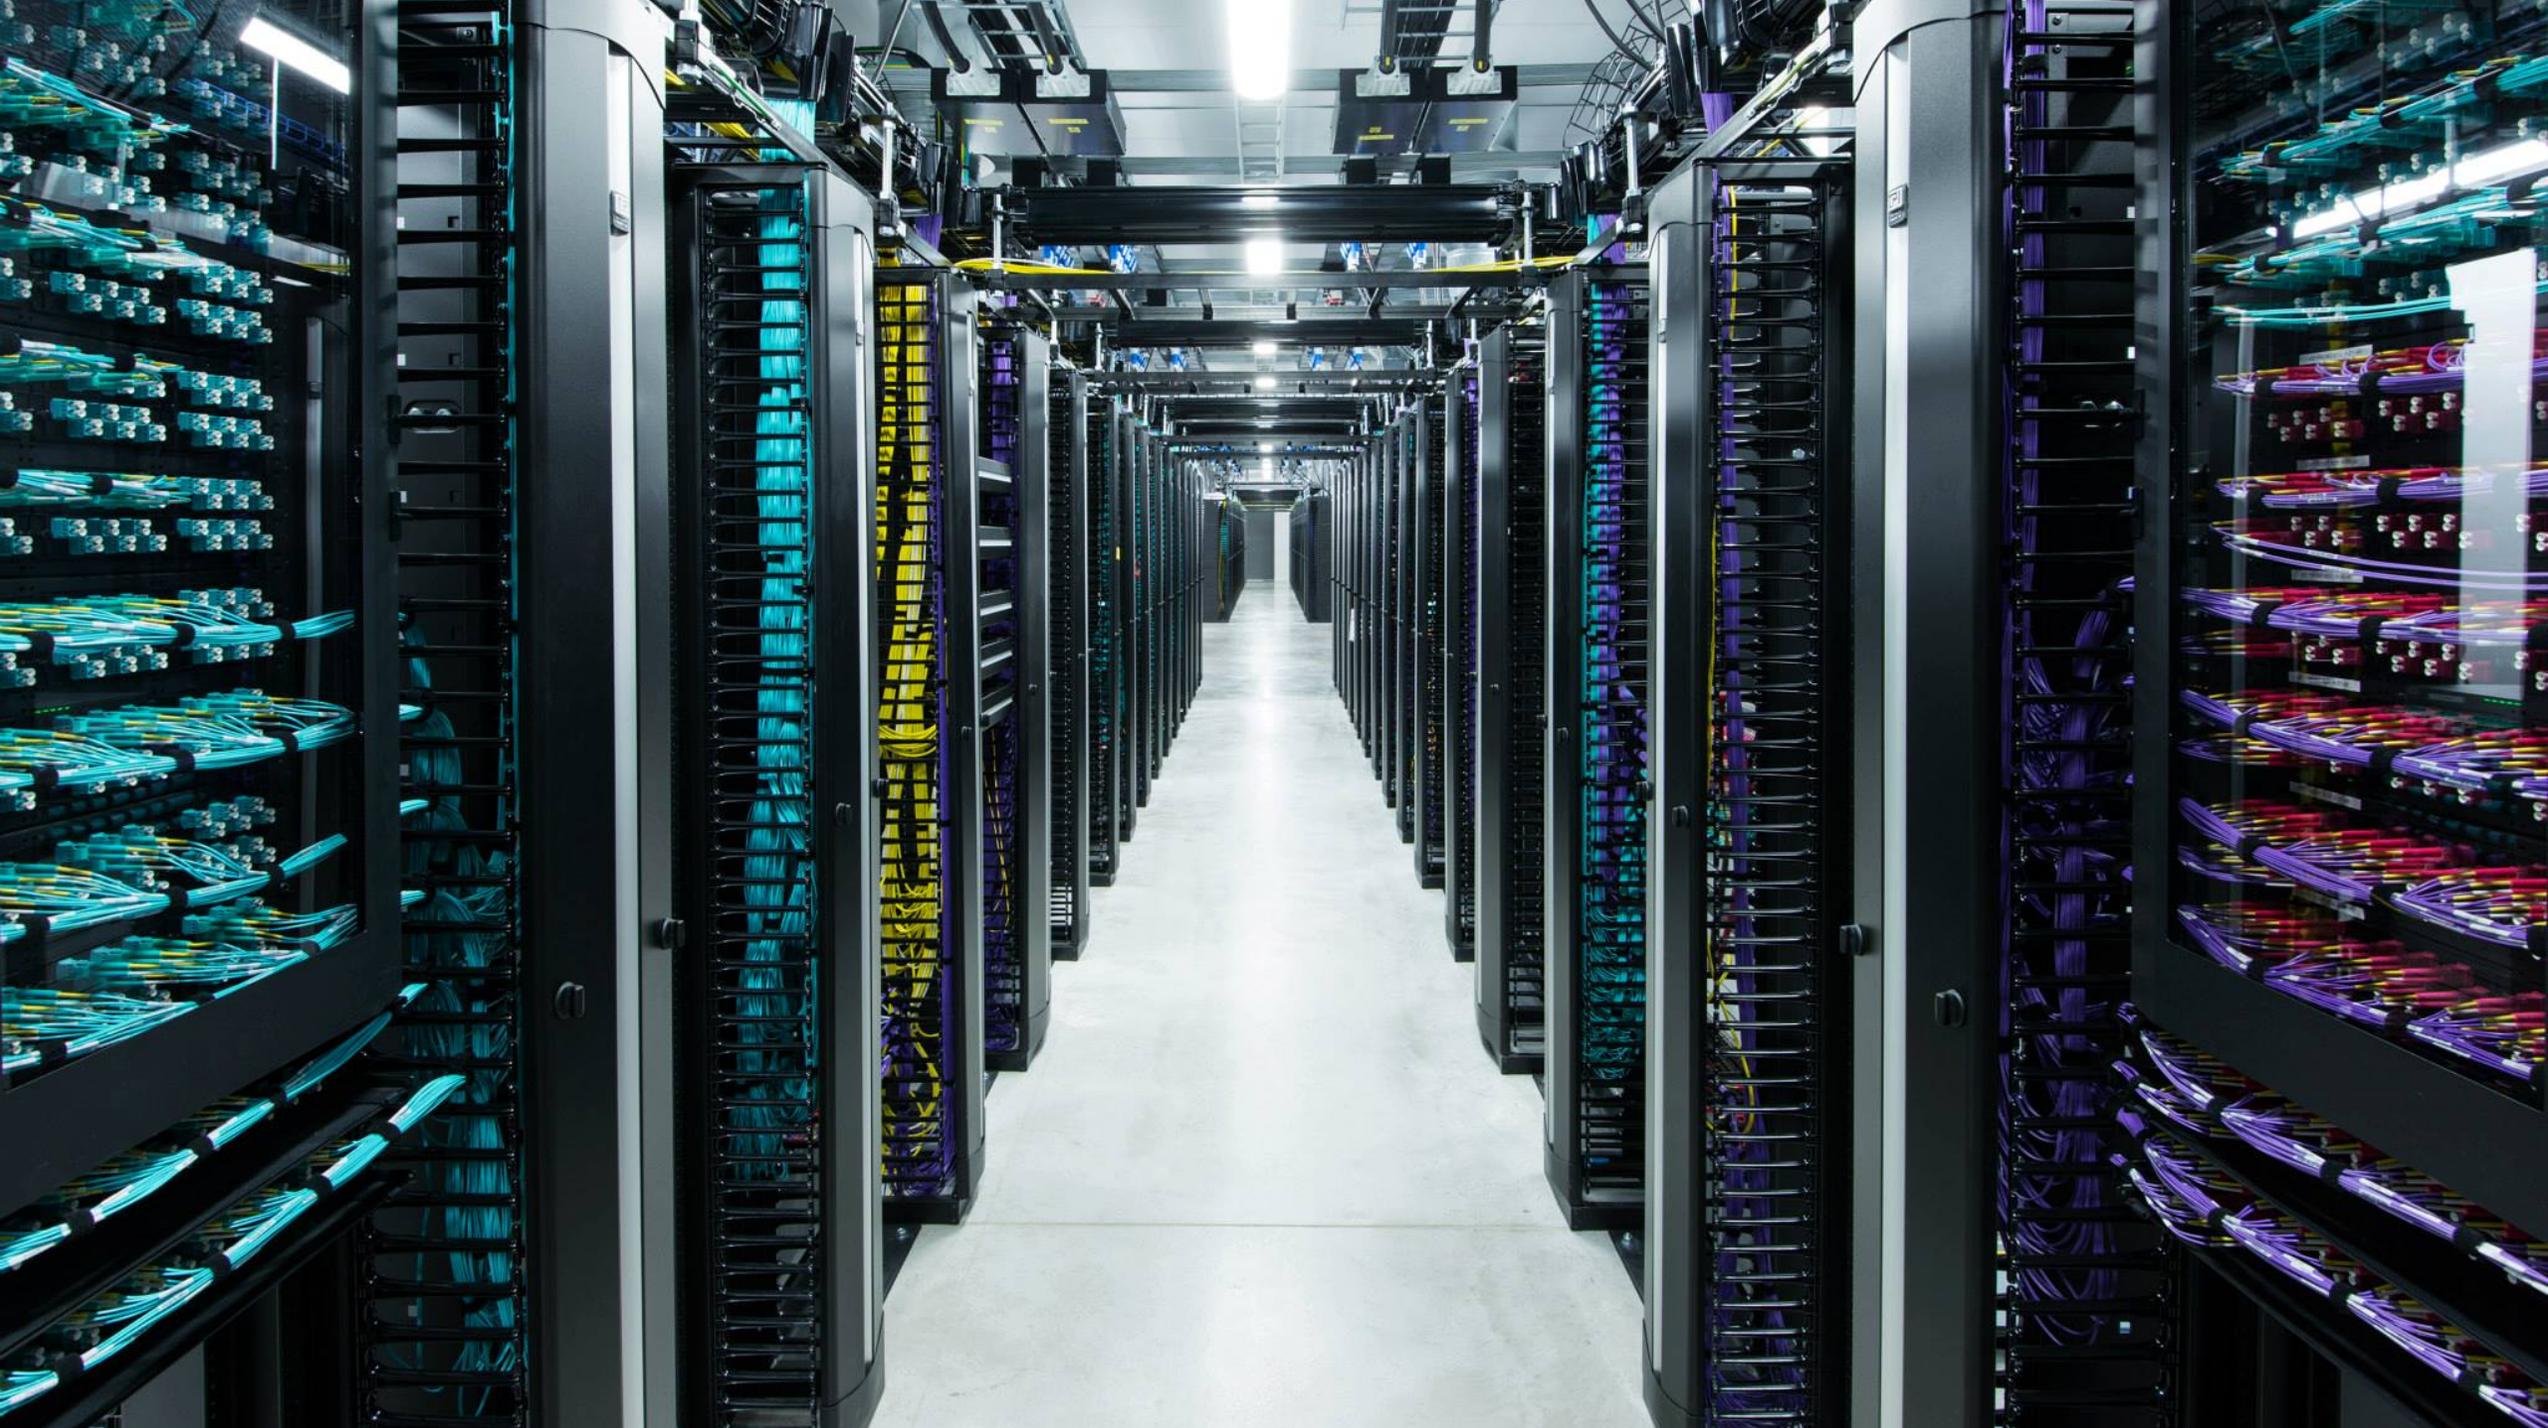




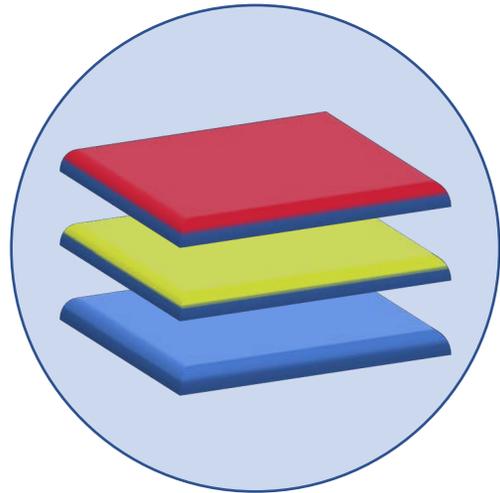
FBOSS Software Team



What can you do?



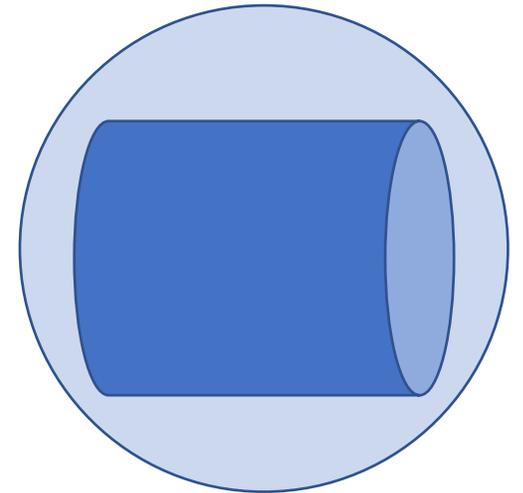
Rules of the Game



LAYER 3



**INDUSTRY
STANDARD**



CAPACITY

No Layer 2

Seriously no Layer2!

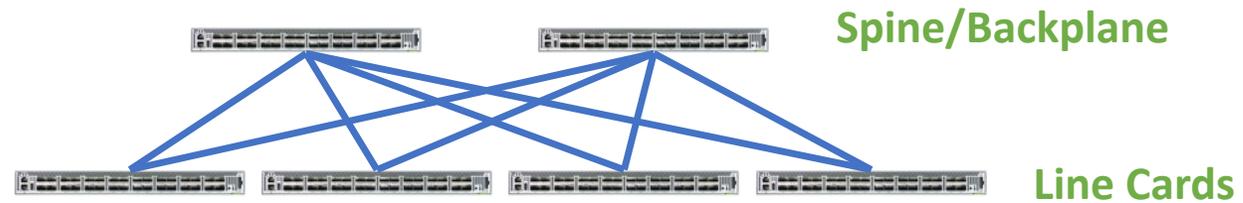
Start with a cheap 32 Port 40G switch



32 x 40G Switch = \$10,000

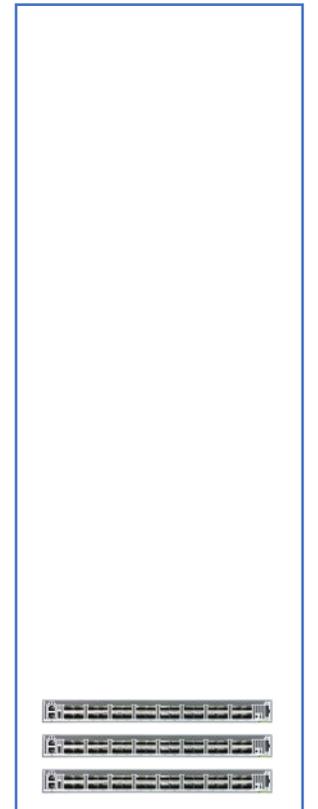
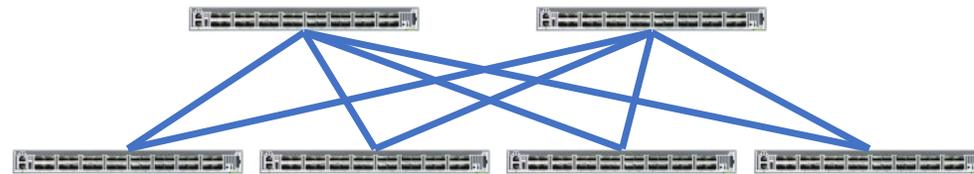
Port cost = \$300

Build a Virtual Chassis

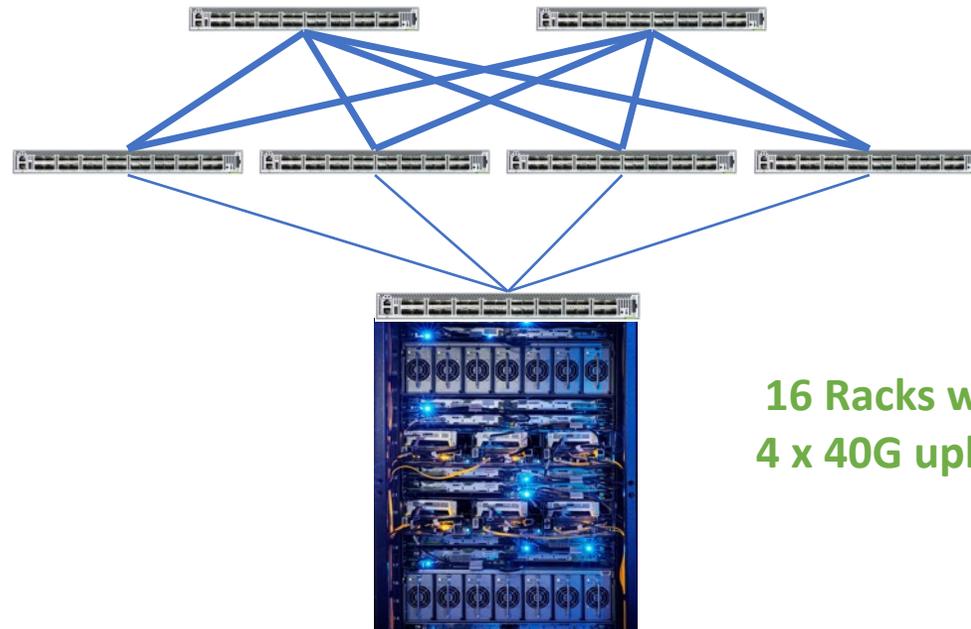


Build a Virtual Chassis

Physical and Logical Redundancy



Attach some Racks

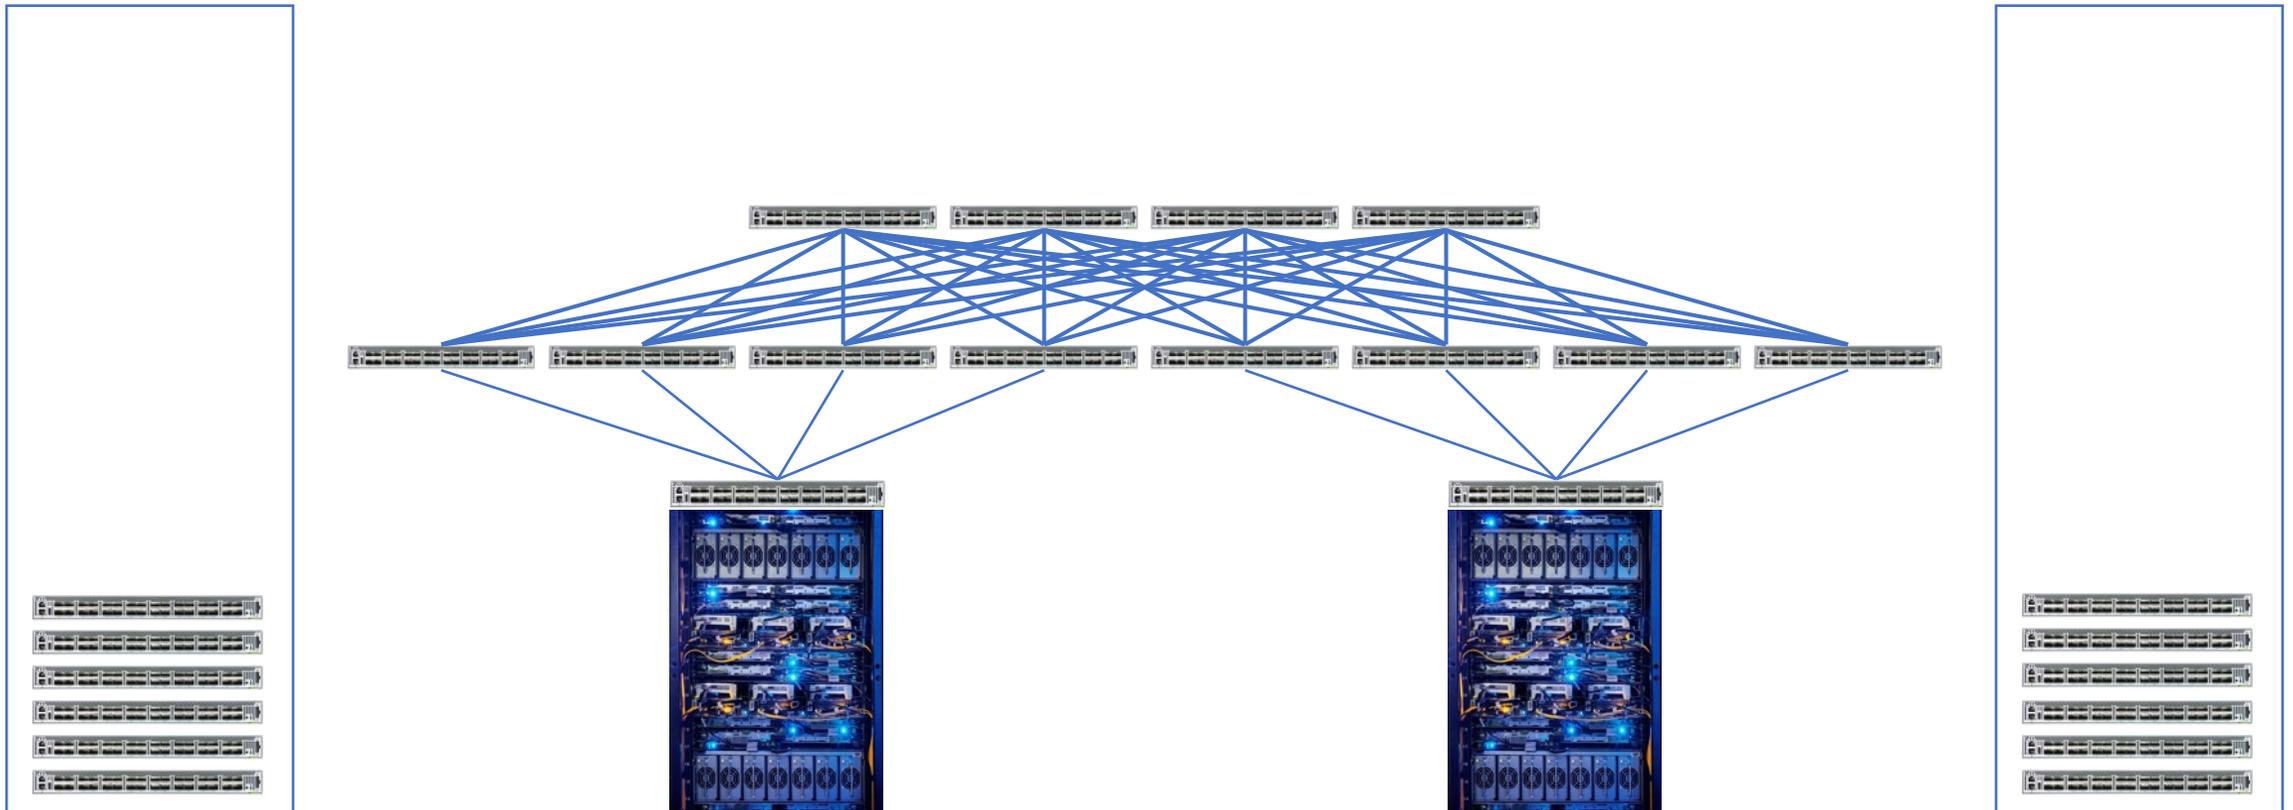


16 Racks with
4 x 40G uplinks



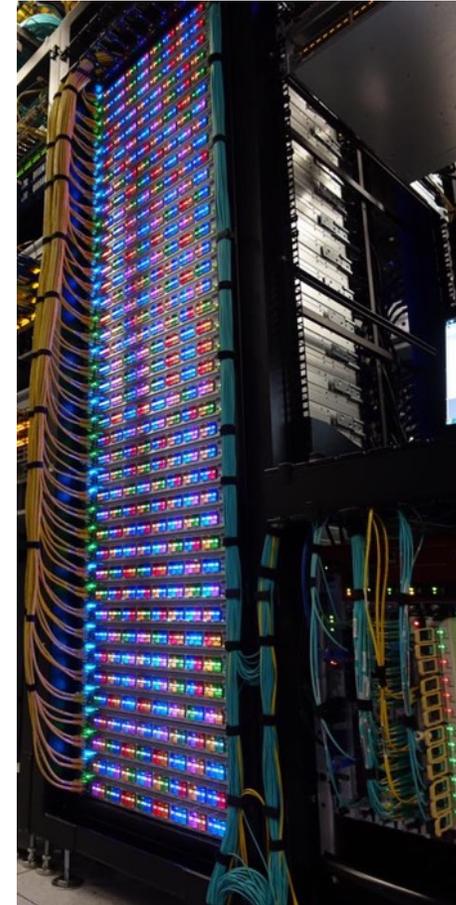
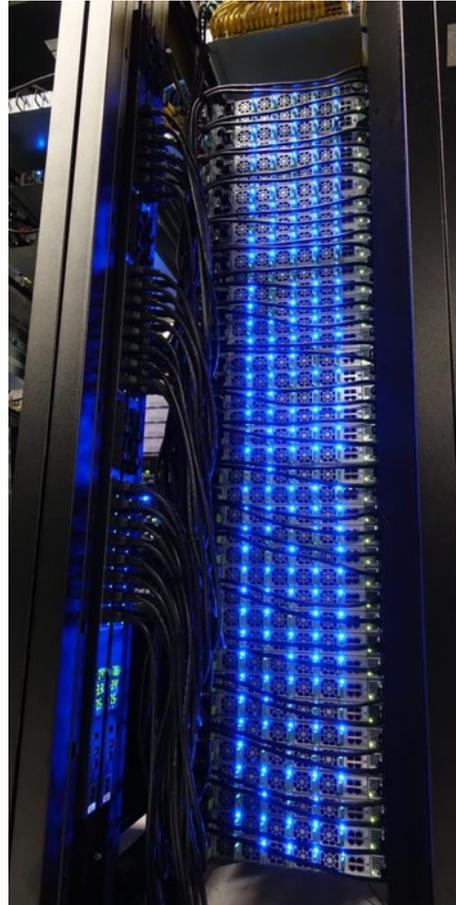
Virtual Chassis Scaling

- No Single Points of failure
- Any individual device failure is just 25% of capacity



How many 32 port switches fit in a Rack?

- 40 Devices x 32 ports = 1280!
- 640 non-oversubscribed rack facing ports



Small versus Big



Fixed Switch



Chassis



- **Smaller Table Sizes**
- **Smaller TCAM**
- **Smaller buffers**
- **No fancy chassis features**

- **Route aggregation**
- **Restrict use of ACLs within dc network**
- **Add bandwidth so buffers empty faster**



ISSU

Provisioning

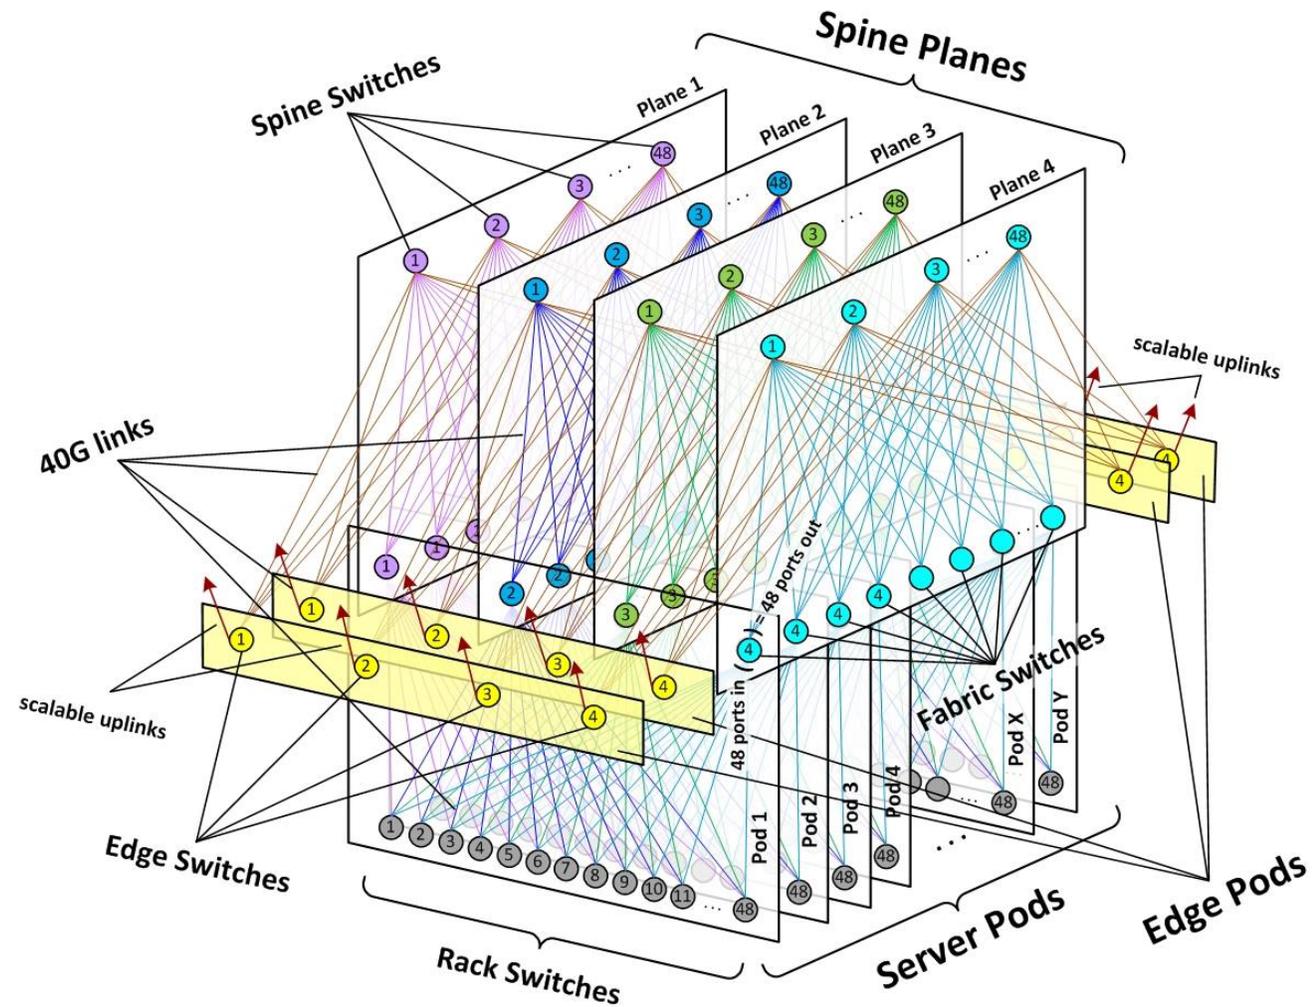
- Standardize switch configuration
- Source control
- Build simple tools to take switches in/out

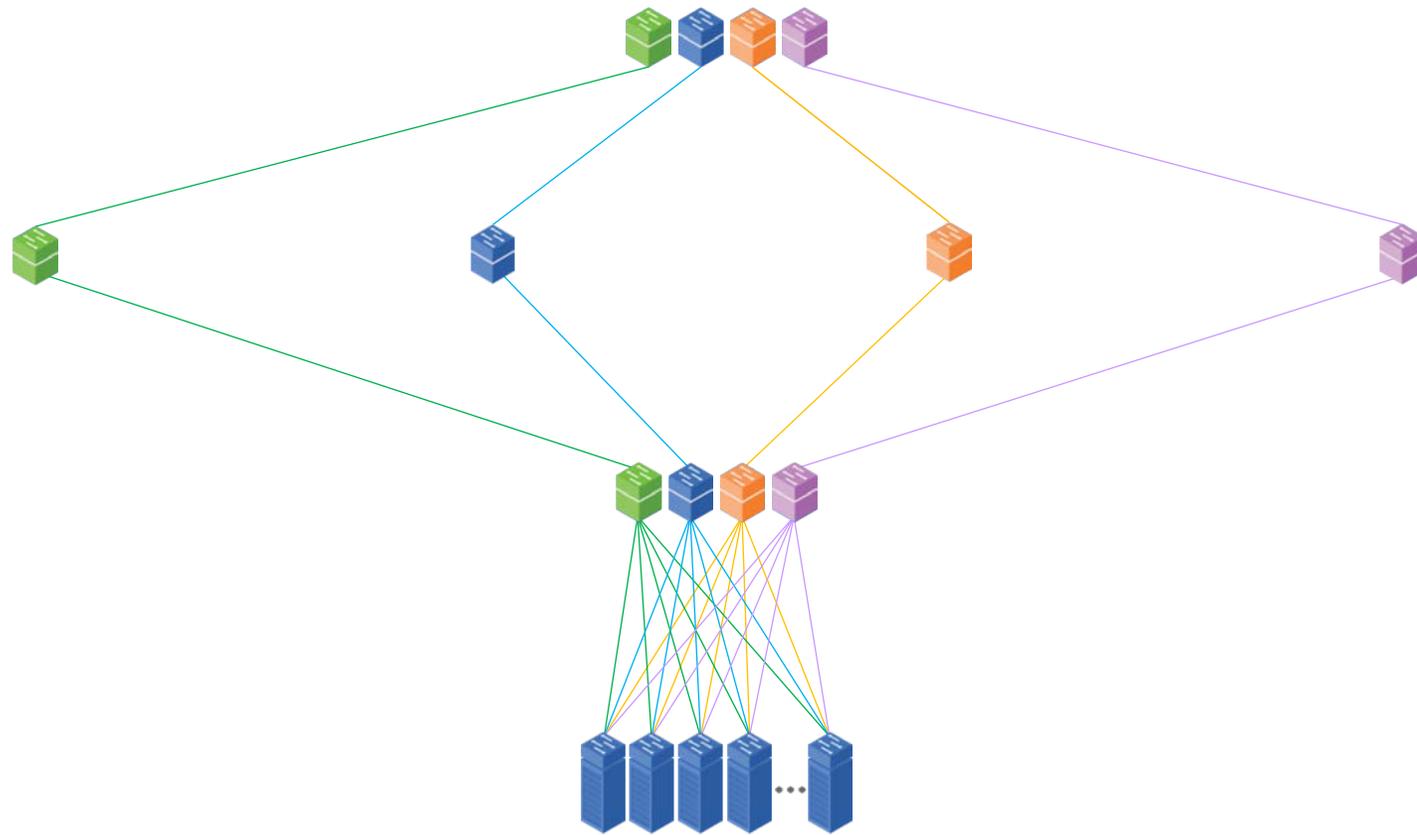
Monitoring

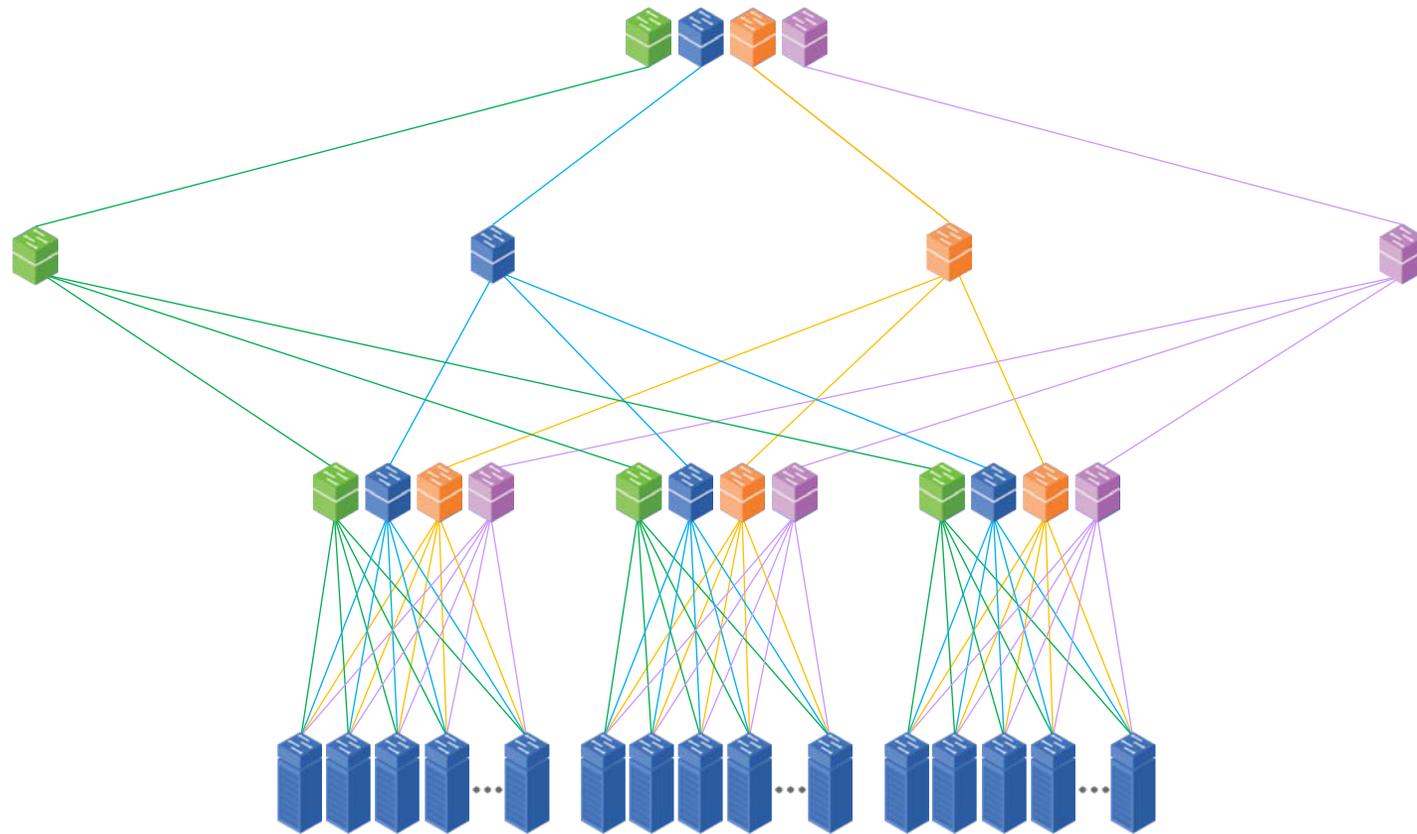
- Central logging of syslog, SNMP
- Ping all the things
- When a switch starts reporting errors take it out

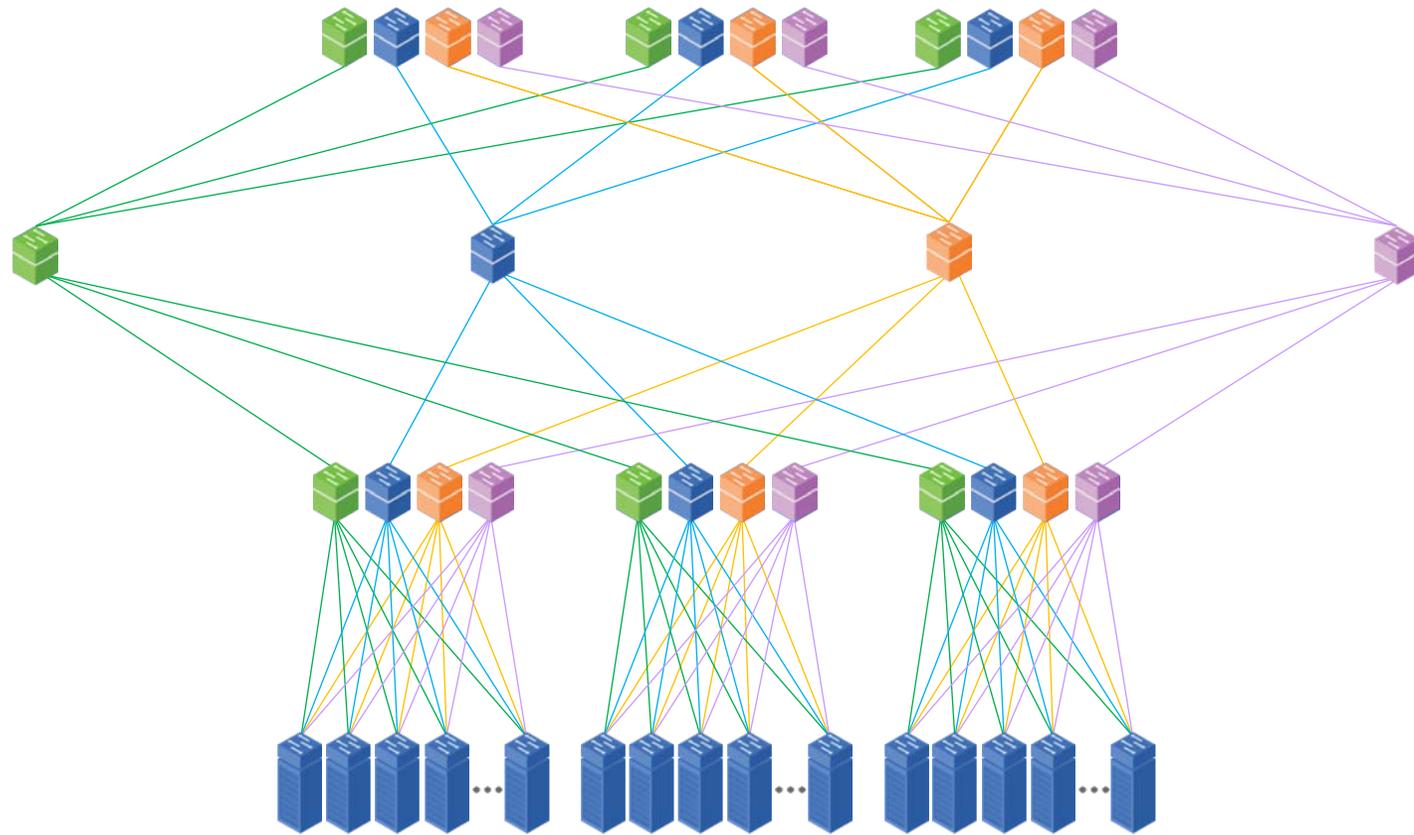
Open your mind

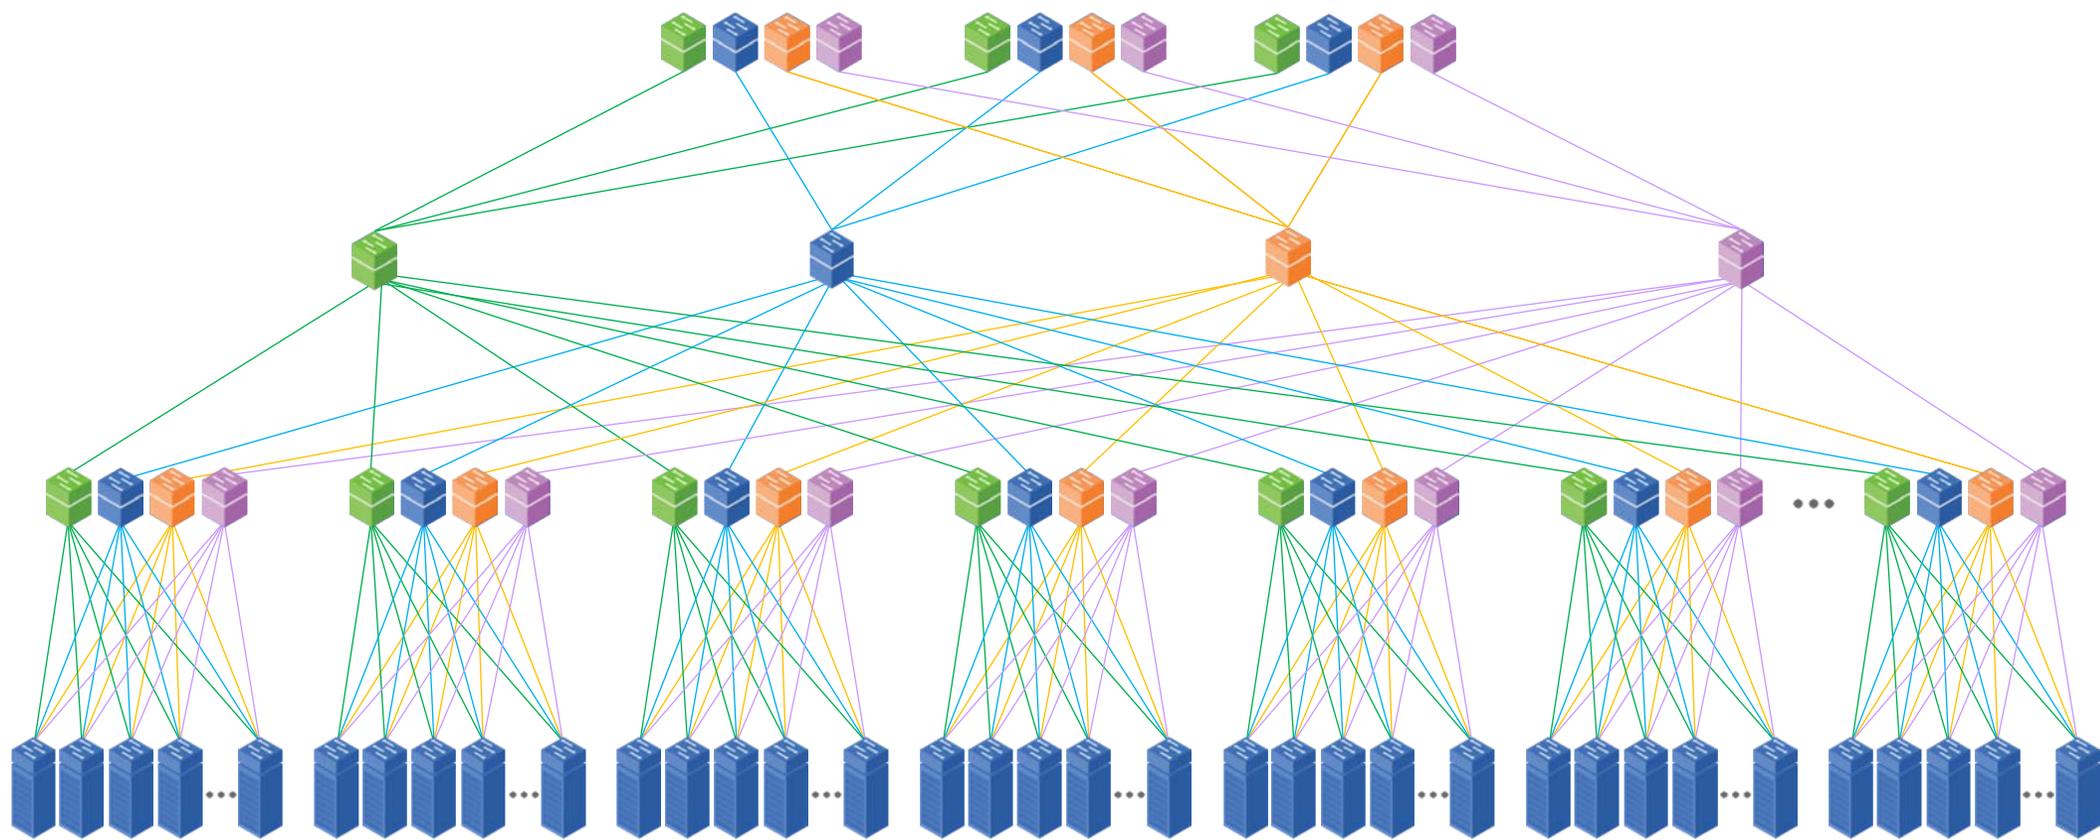
Datacenter Network is a giant Virtual Chassis!

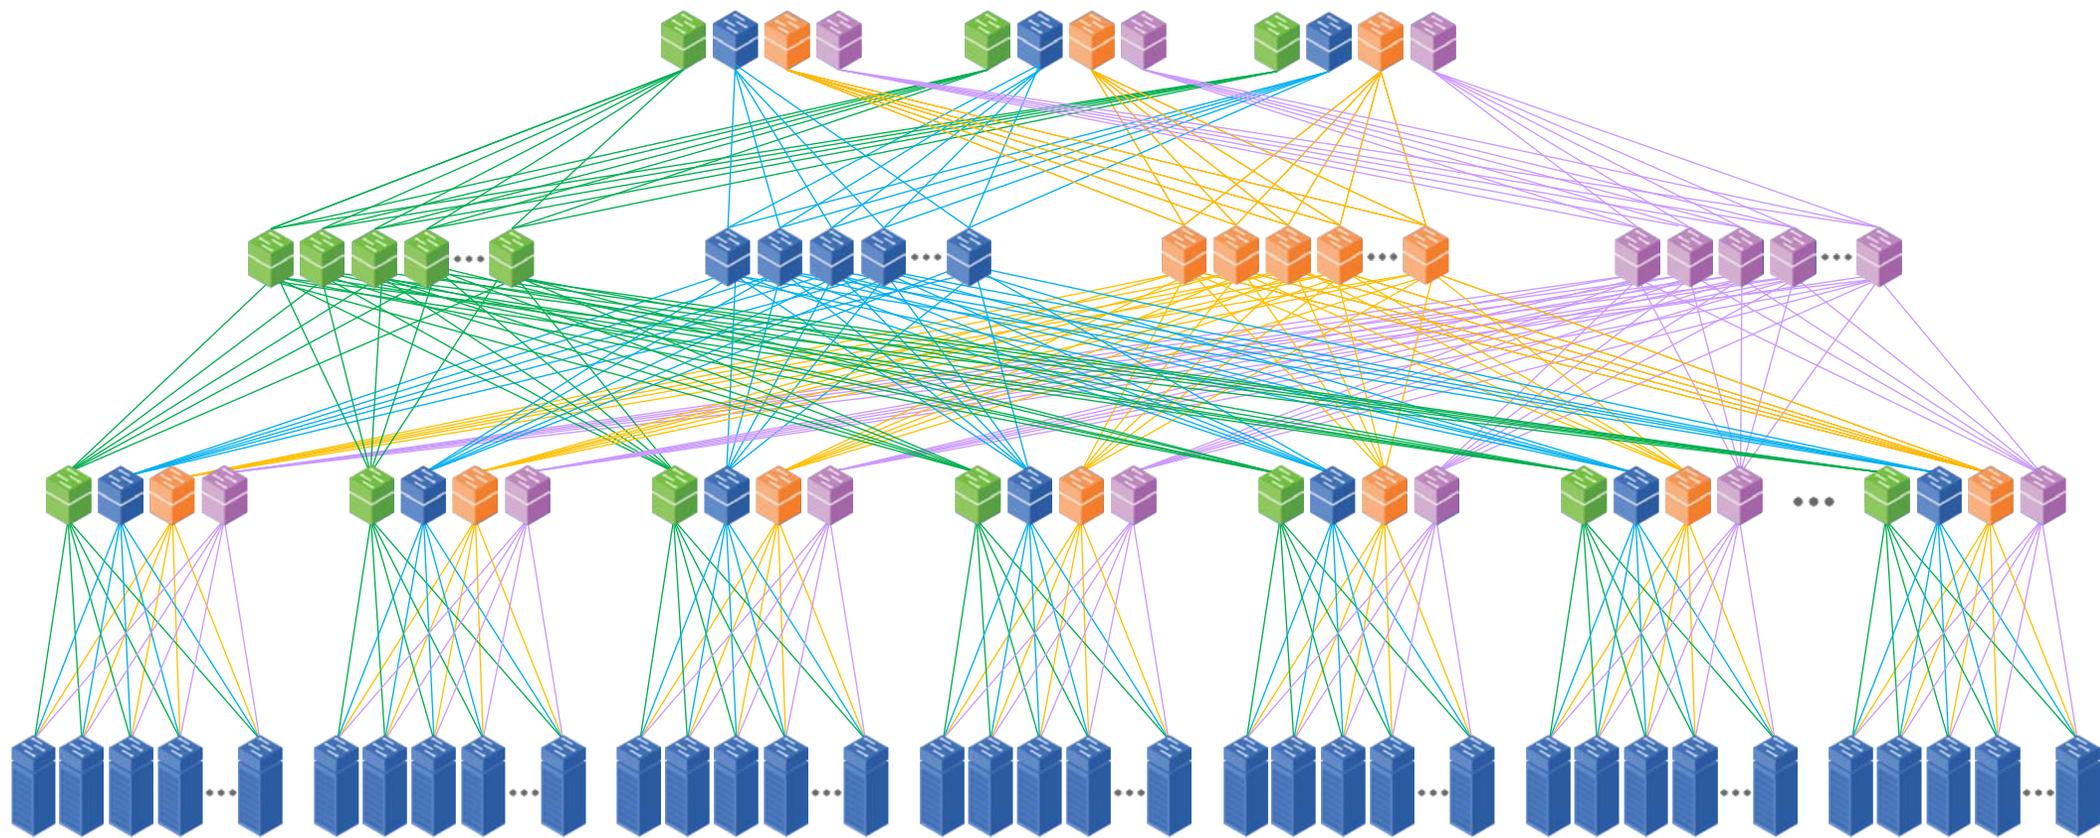


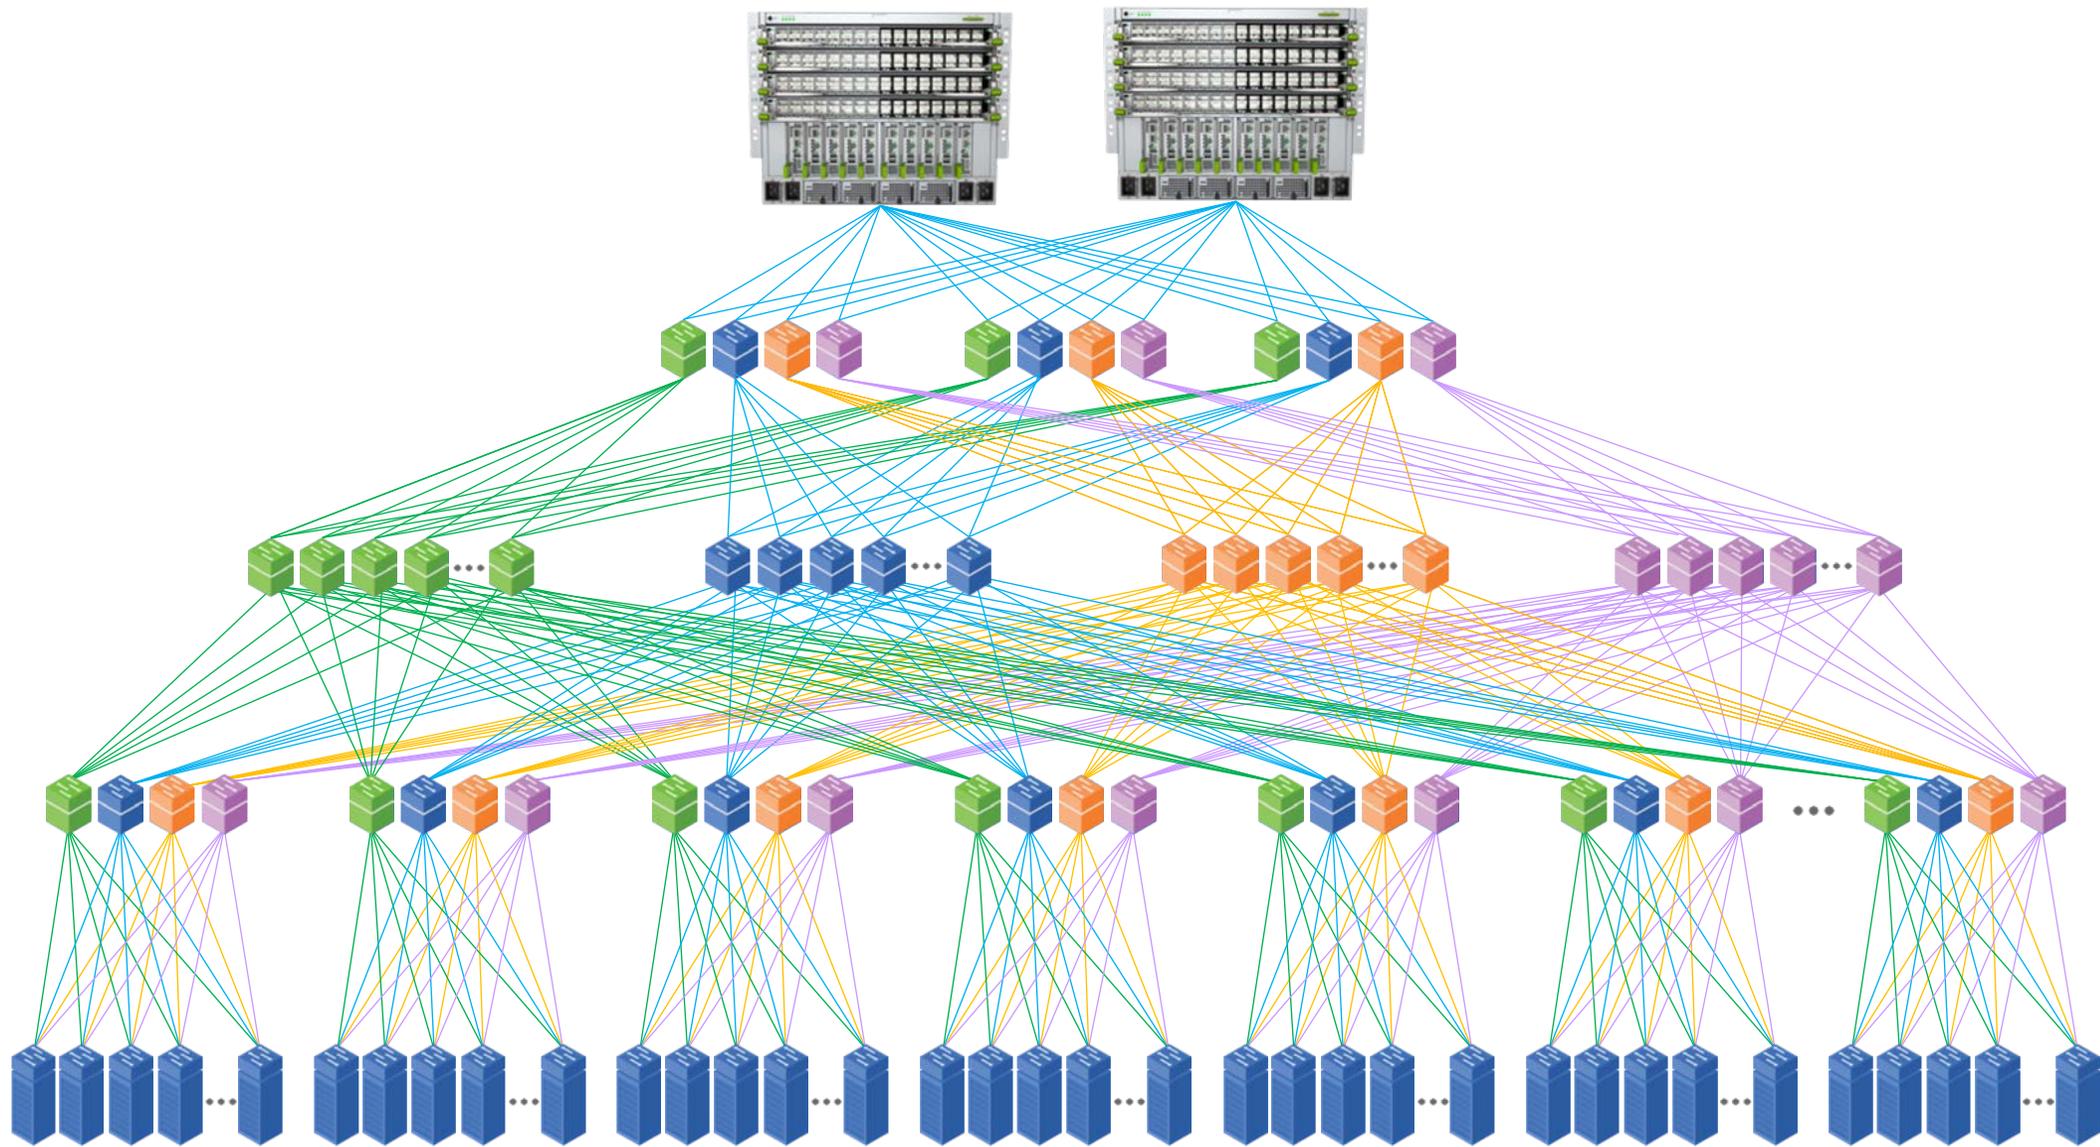




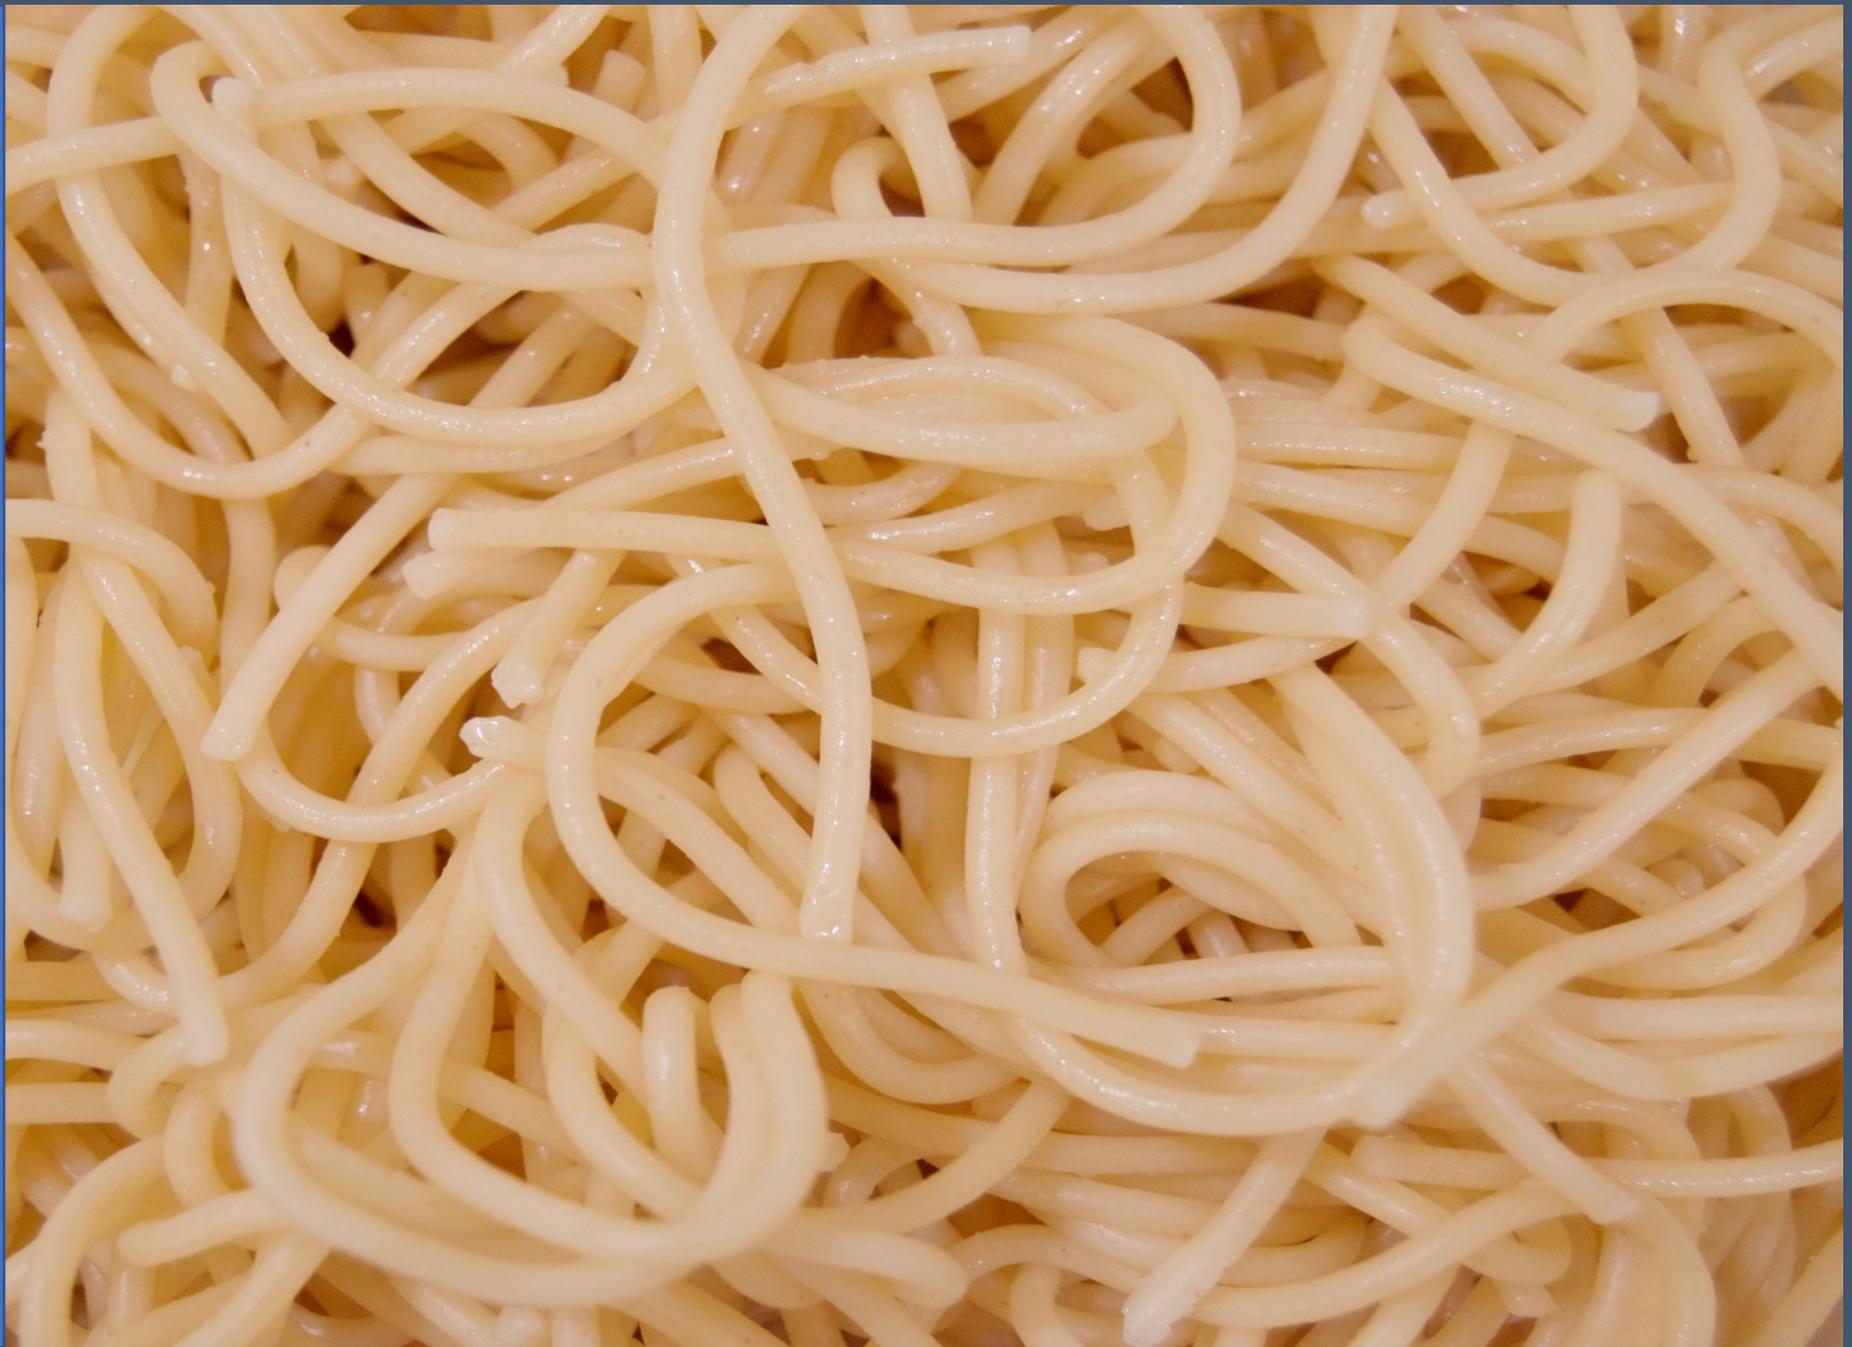


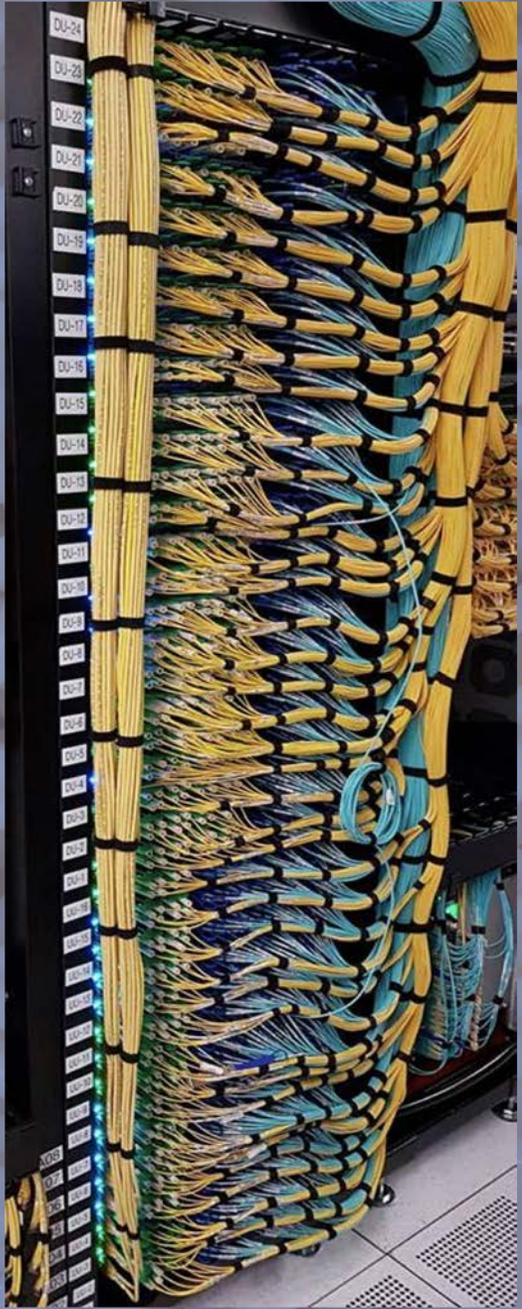






Cabling





32 data cables
1 oob
1 mgmt
+ 2 power

36 cables x 40 = 1,440 cables!

- Color coding
- Take time to cable right
- Cable selection - AOCs, DACs



What's in it for you?



Thank You

facebook