

The Future of Datacenter Networking

Presenter: Ana Radovanovic
(anaradovanovic@google.com)

Google Network Infrastructure Teams

Internet Growth Rates




**Internet Observatory
Report**
CAGR: 44%

MINT Study
CAGR: 50%

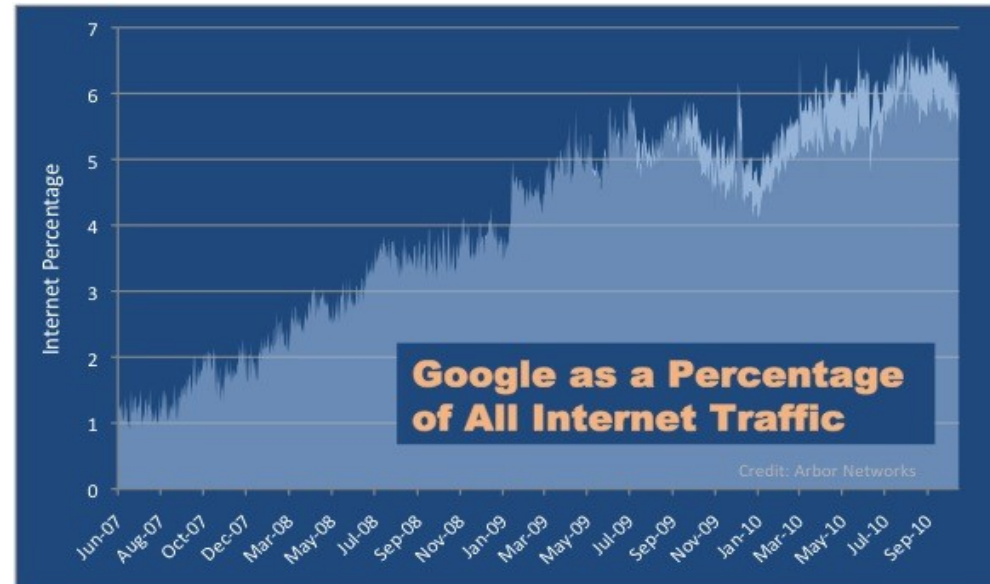
<http://www.dtc.umn.edu/mints/home.php>

Cisco Study
CAGR: 34%

http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html



(a) Example AGR Calculation

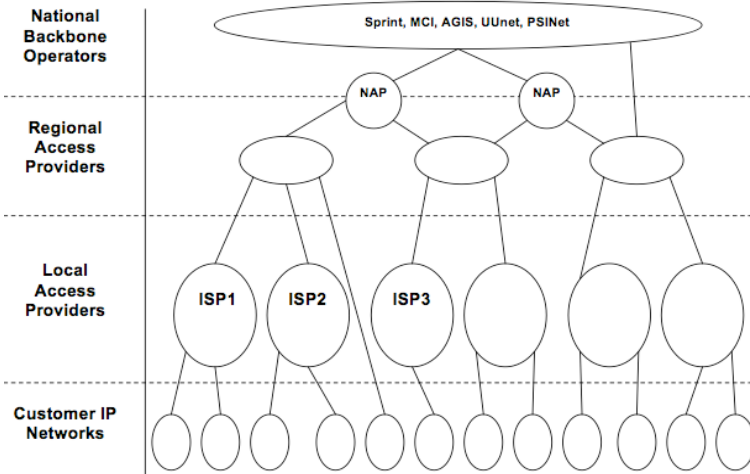


**Global Internet Traffic is growing at 34%-50%
year-over-year rate**

Internet Topology Evolution

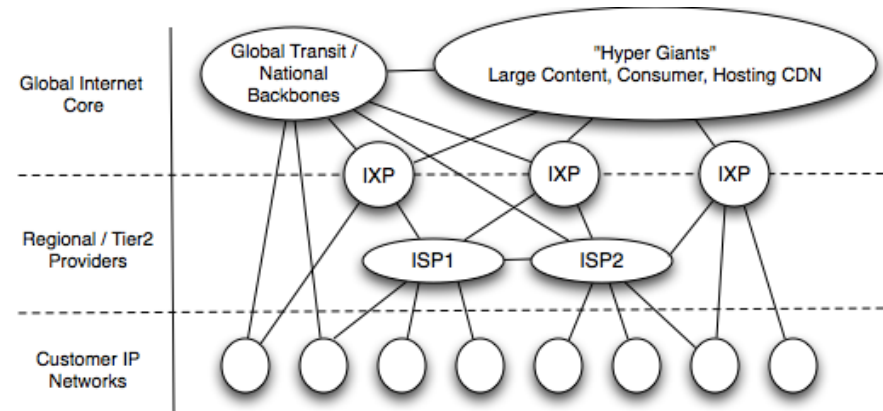


Textbook Internet 1995-2007




Hierarchical, Tier 1 Focused

~ Internet Today



Content 'Hyper Giants': direct connection of content and consumer



Consolidated Computing, Many UIs, Many Apps, Many Locations

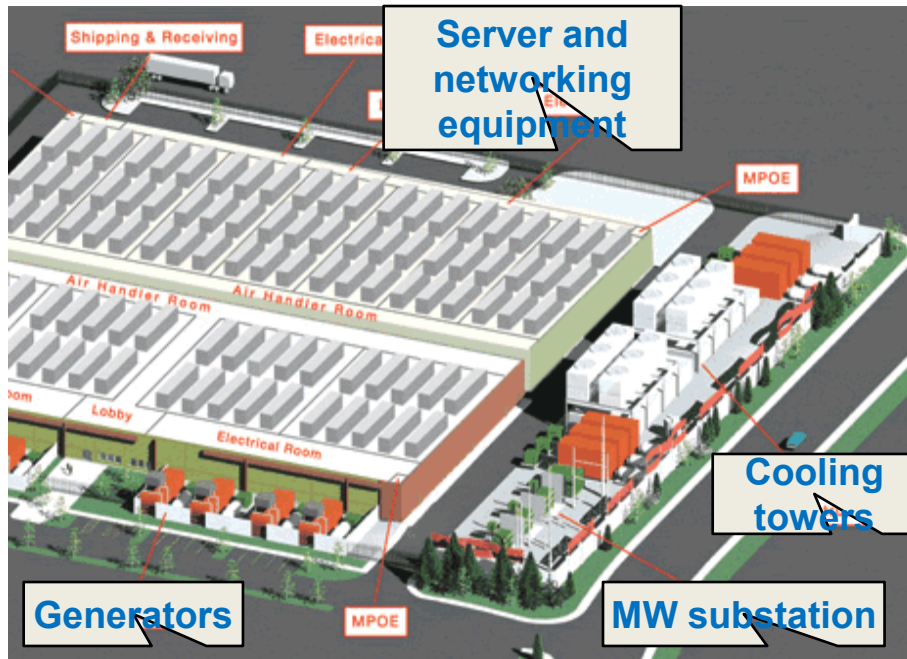




Warehouse Scale Computer

The core of Google's Infrastructure

Warehouse Scale Computer -- Overview




- Collection of servers
 - Services such as Search execute at a **scale** far beyond single machines or racks -- require no less than clusters of 100s or 1000s of machines
- Machines, Network, and Software all working in concert to provide Internet scale services – the data center *is* the computer

WSC Building Blocks



A datacenter contains 1 or more *clusters*,
and has a *network* and a *power topology*



rack: 40-80 machines
+ Ethernet switch


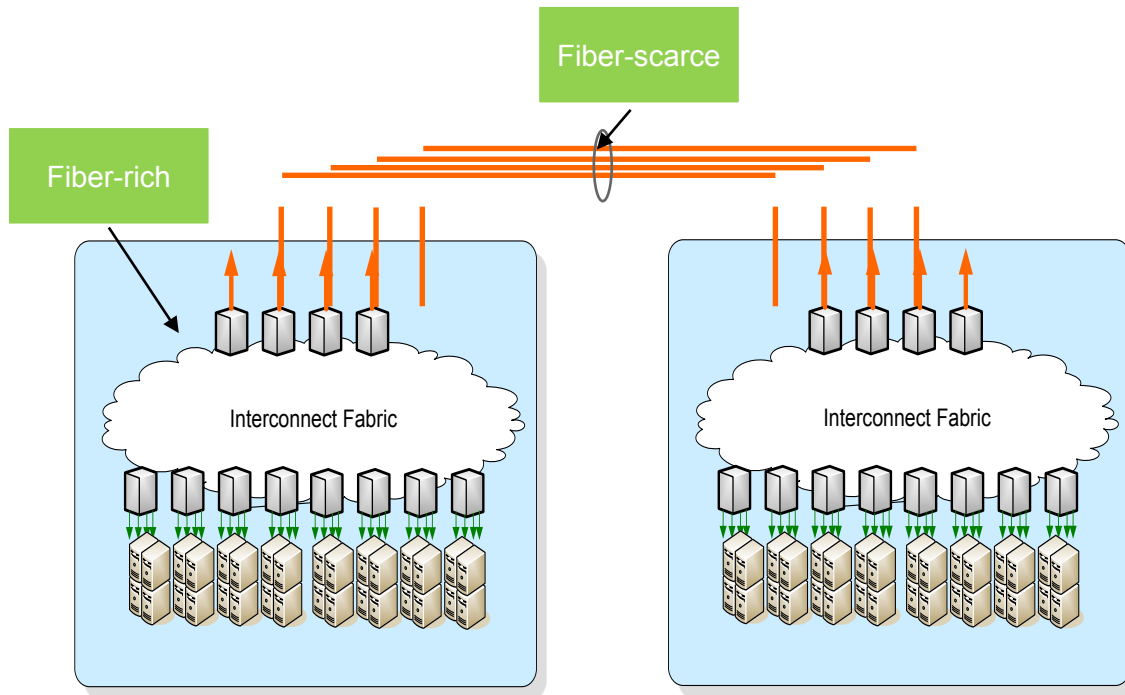
Warehouse Scale Computing Characteristics

- Relatively homogeneous machine, network and systems software platform
- Common systems management layer
- Massive **Scale**, driven by:
 - User Base (*100s* of millions of users globally)
 - Data Set Size and Growth (Search --Web corpus/Youtube -- Video corpus)
 - Introduction of Novel Features (Instant Search, HD Video)
- Efficiency, driven by **Scale**
 - This relentless demand for more computing capability makes *cost efficiency* a primary metric in WSC design

WSC Network Challenges/Constraints




Traditional DC networking components systems and protocols impose constraints that are counter to the goals of WSC. New Solutions are required to meet our requirements



Traditional Distributed Control/Management

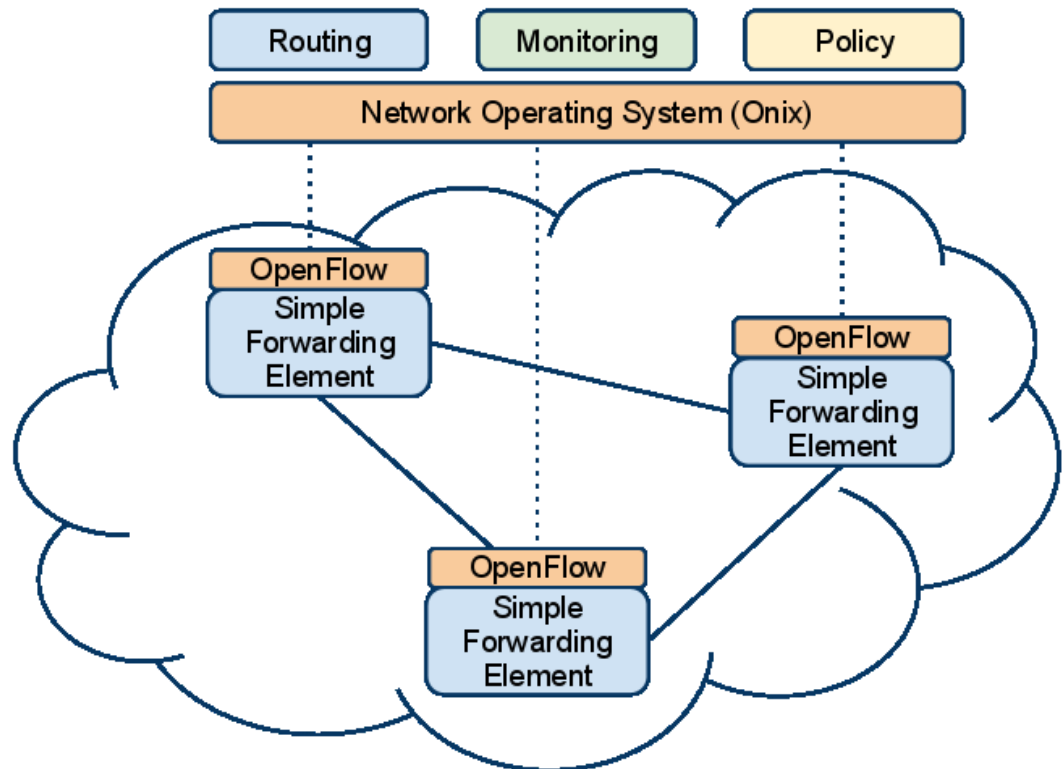
- Difficult to implement and maintain unified configuration and policy
 - Many systems, many configurations
 - Prone to Human error (a leading cause of outages/unavailability)
- Little Service/Application Awareness
- Complex, proprietary protocols
 - Difficult to change, deters innovation
- Proprietary management systems
 - Don't scale themselves
 - May not interoperate
- Programmability actively discouraged
- *Scales poorly*



Software Controlled Networking



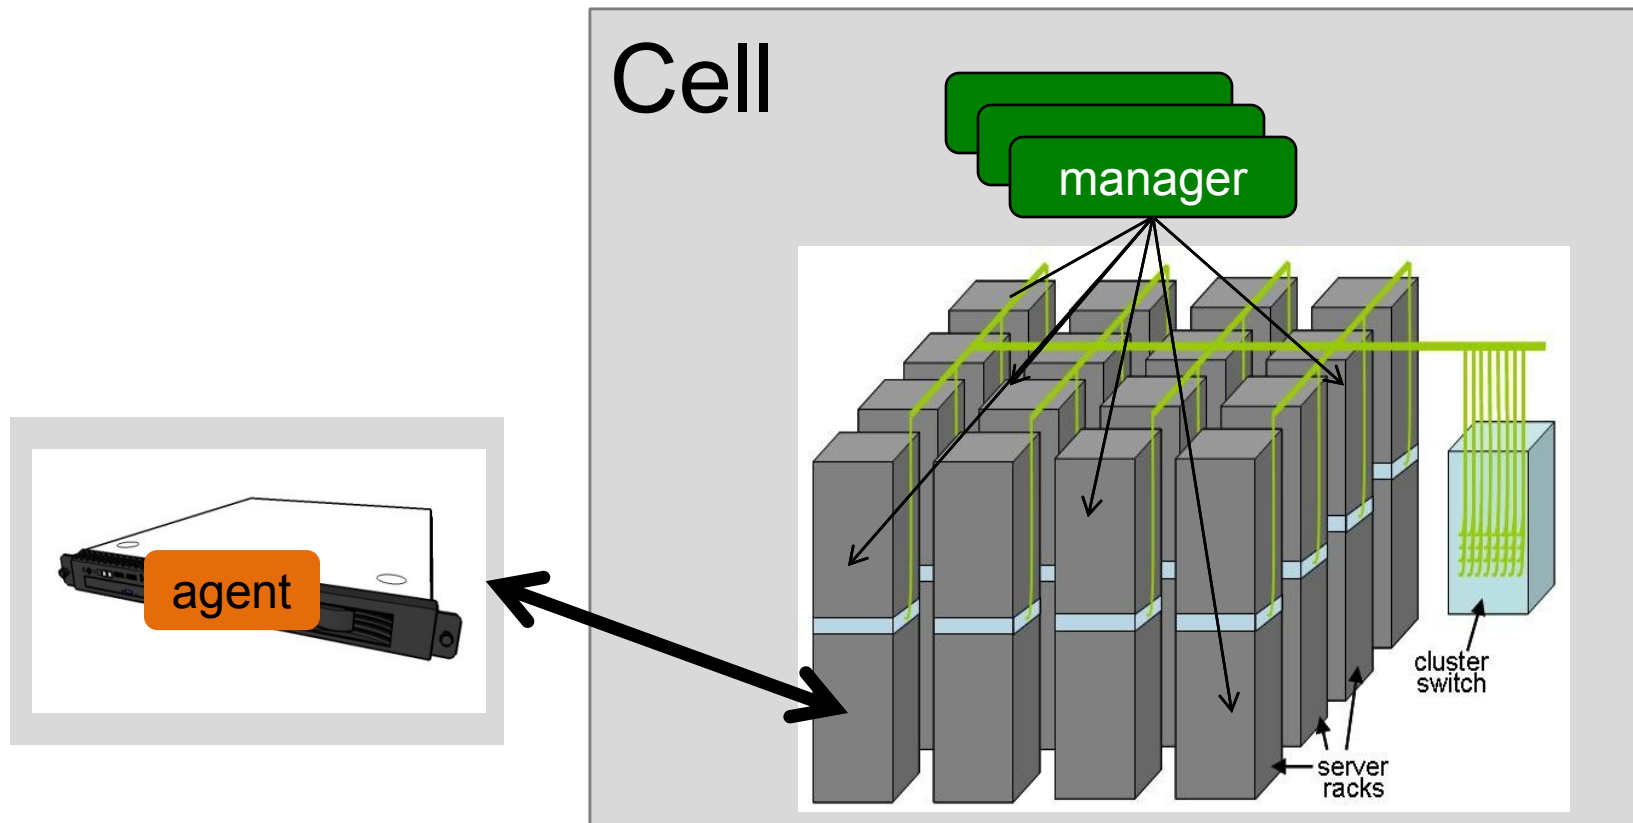
- Centralized, 'Programmable' network model
 - Simple to maintain cogent, unified policy, configuration across many and diverse network elements
 - Service/Application awareness and integration, improved alignment of Biz priorities and resource allocation
 - Supports rapid innovation
- Scales well
 - Including WSC footprints
- Existing solutions
 - OpenFlow, Onix



Cluster management: what is it?




- Each cell has a (replicated) central *manager*
- Each machine has a local *agent*



Making It All Happen



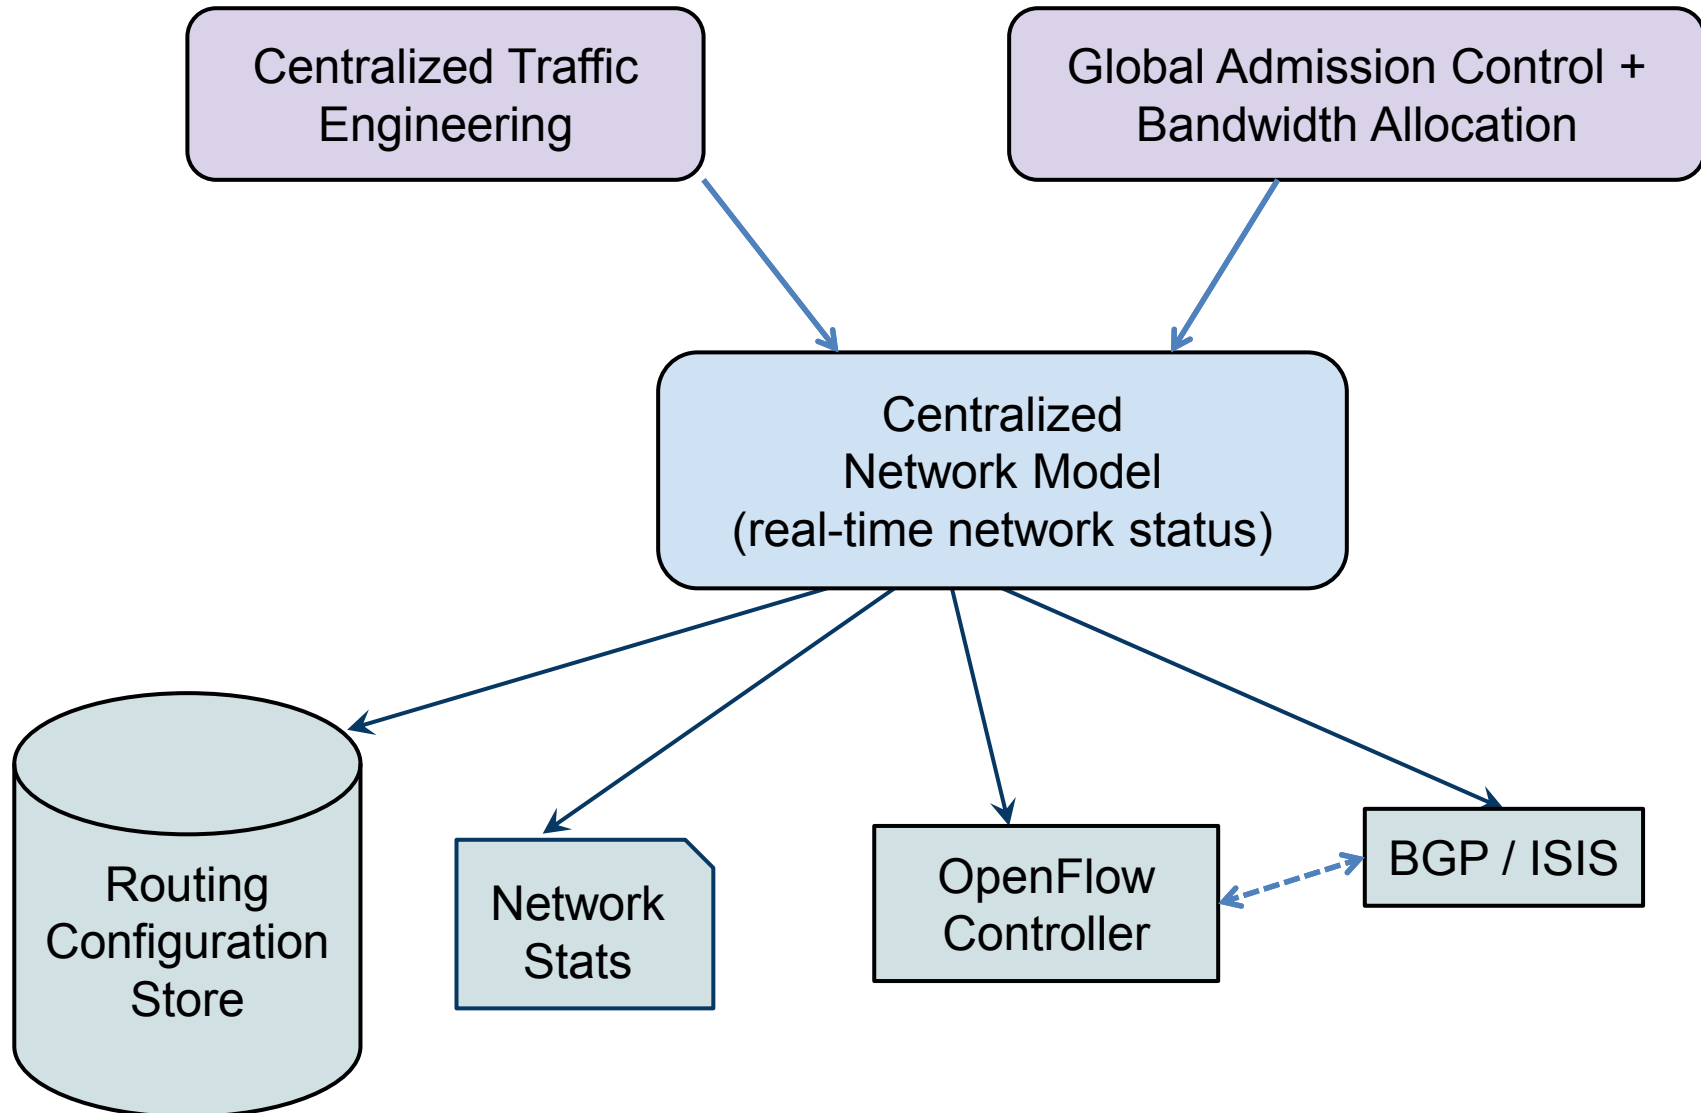
- Physical layer: WDM, SMF and OCS (cheaper!)



10x10G CFP

- Centralized control:
 - Better network utilization with global picture
 - Converges faster to target optimum
 - Allows more control and specifying intent
 - Can mirror production event streams for testing

Making It All Happen



- Rise of cloud and content
- Scaling and efficiency are the keys
- Solutions:
 - Scale-out
 - Low cost WDM interconnect and high-radix OCS
 - Software Controlled Networks



Thank You
