deepfield

security | performance | control

The Operational Intelligence Company

The New Internet

Craig Labovitz
labovit@deepfield.net

(Stefan Meinders, stefan@deepfield.net)

The New Internet Data Sources

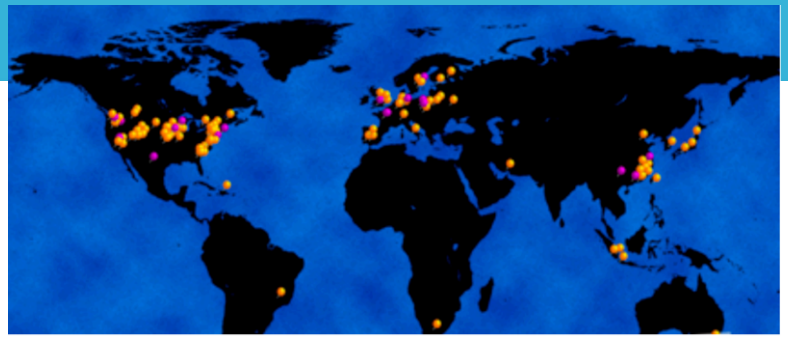

- Large scale non-commercial, research study of intra and inter domain traffic with focus on changing content engineering and interconnection

Sources

- Cross ISP
  - Cable Labs / Princeton / Deepfield
  - Seven publicly announced participants
  - >50% US Internet consumers
- Internet Observatory
  - Research collaboration (Deepfield, University of Michigan, Merit Networks)
  - Includes fixed-line, CDN and Content
  - Statistical significant consumer traffic (North America)
  - <10% Europe

Cross ISP

http://arxiv.org/pdf/1603.03656.pdf

Participants

bright house NETWORKS · COMCAST · COX · Mediacom · MIDCO · sudden link · Time Warner Cable

CITP · CableLabs · deepfield

Data Collection

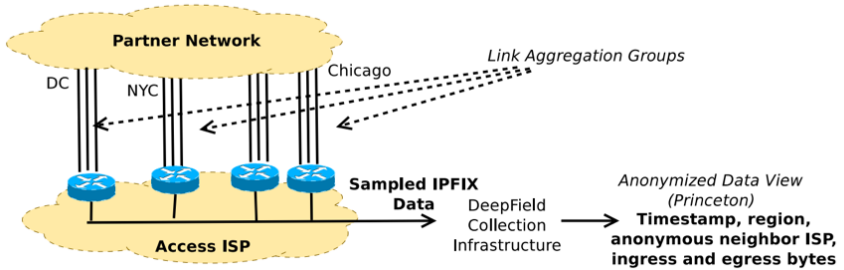

- One year collecting **SNMP** (capacity), **BGP** (aspaths), and **Flow** from the majority of peering interfaces across seven US ISPs. Additional data collection from internal CDN and middle boxes. Fine grain (unpublished) data on streaming bit rate per peer, service, interface and CDN across 50 OTT services

Internet Observatory

http://conferences.sigcomm.org/sigcomm/2010/slides/S3Labovitz.pdf

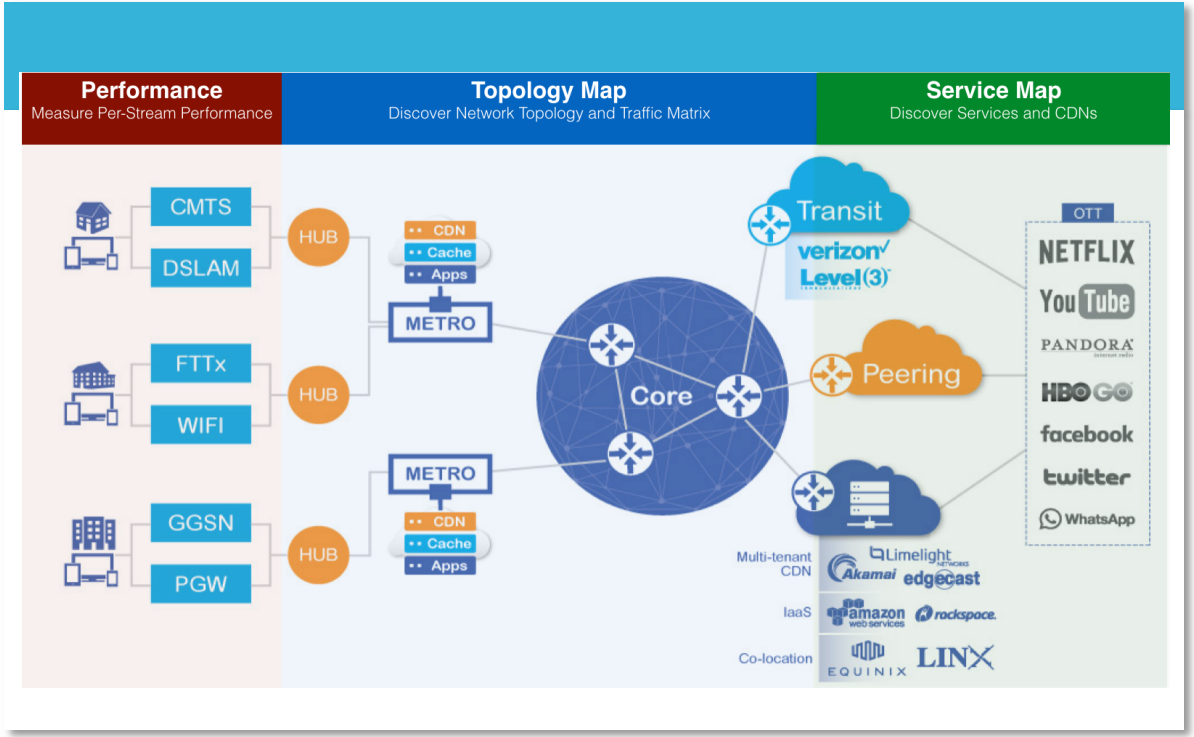

Data

- Research collaboration between UM, Merit, Deepfield and participating providers
- Aggregated data from routers, DNS servers, middle boxes and security (DDoS) events
- Data on CDN, service, tonnage, stream, quality and cyber supply chain
- Daily crawling / nmap of IPv4 and IPv6 address space

Cross ISP Capacity

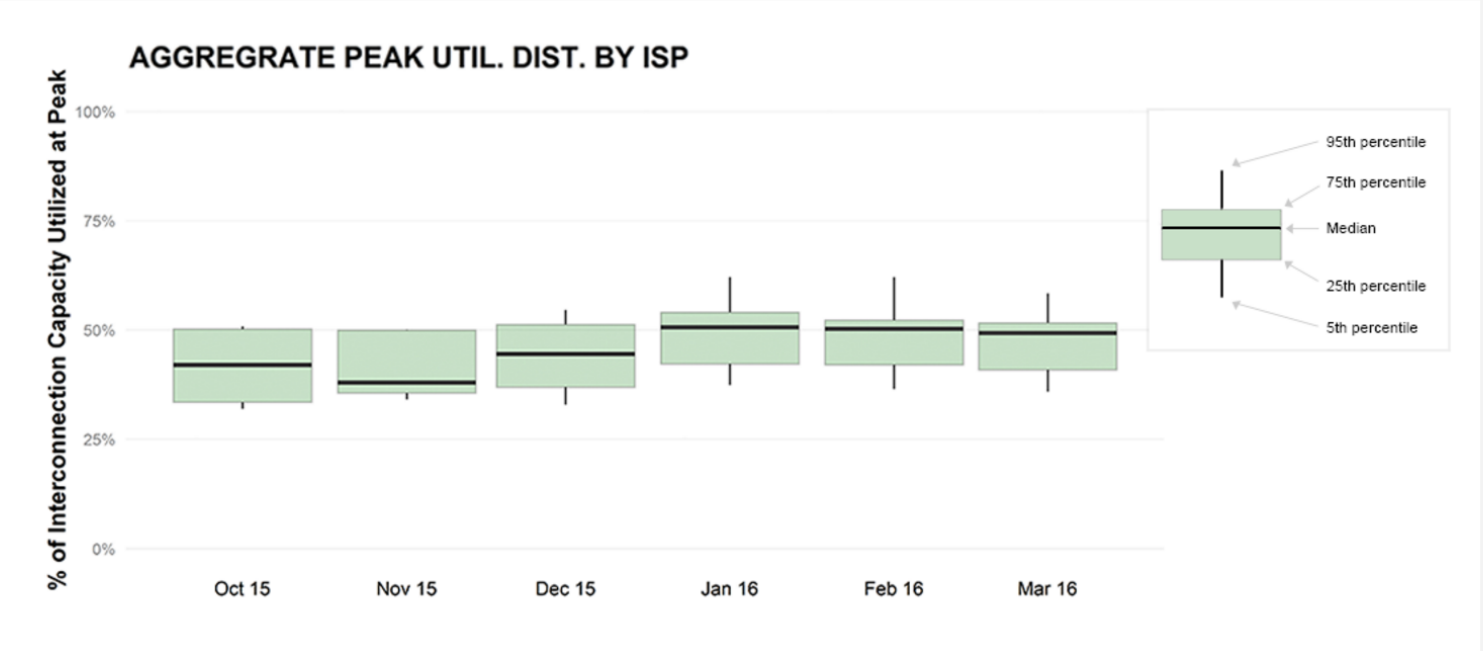

Across all interconnects in the data set, capacity is roughly half utilized during peak periods. During the time period available, capacity has grown nearly 25 percent cumulatively (or, between approximately 2 percent and 5 percent per month), while usage is also growing, though at a more variable rate. Variability in usage growth on a month-to-month basis is likely influenced by a number of factors, including higher usage of Internet video over the winter holidays.

Cross ISP Capacity

CENTER FOR INFORMATION TECHNOLOGY POLICY
AT PRINCETON UNIVERSITY

Interconnection Measurement Project

An Unprecedented Look into Utilization at Internet Interconnection Points

MARCH 11, 2016 BY NICK FEAMSTER

Cross ISP Capacity

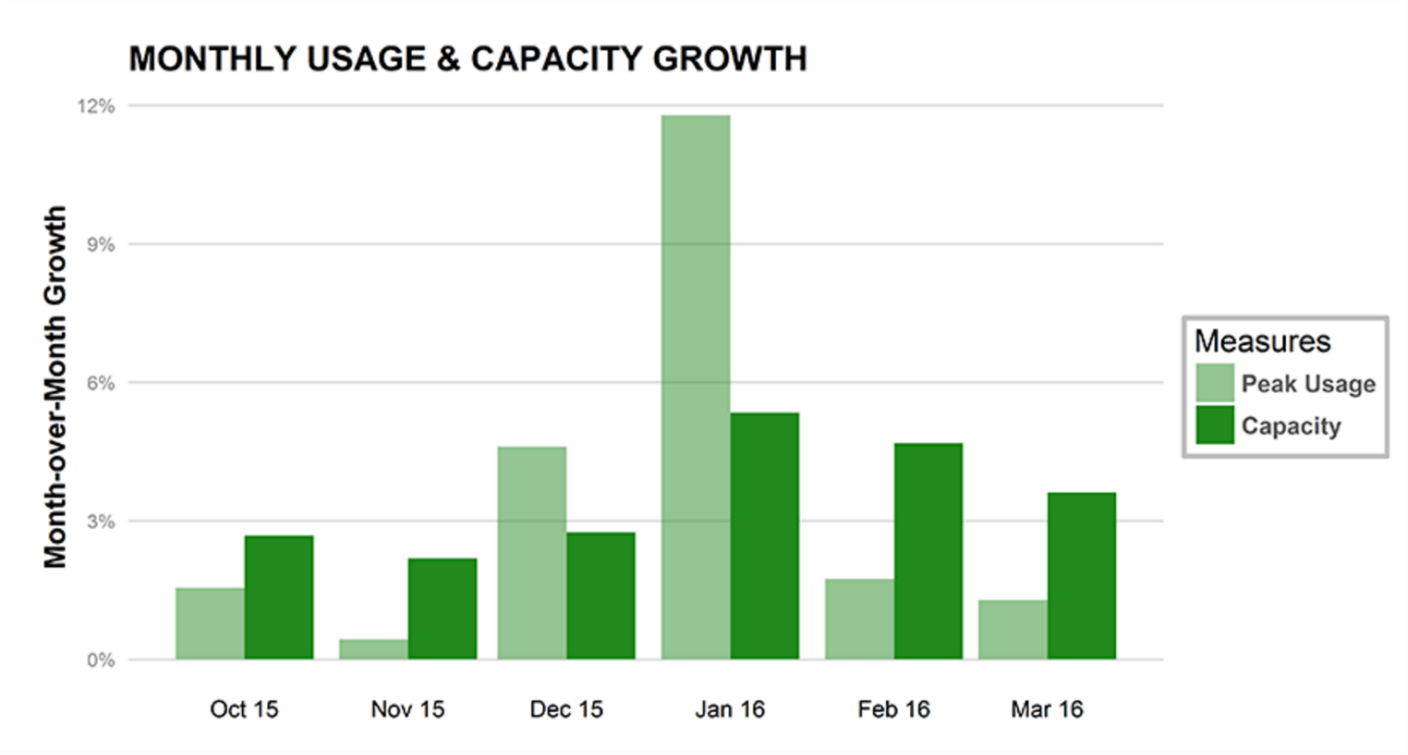

Internet Trends

2007 / 2009

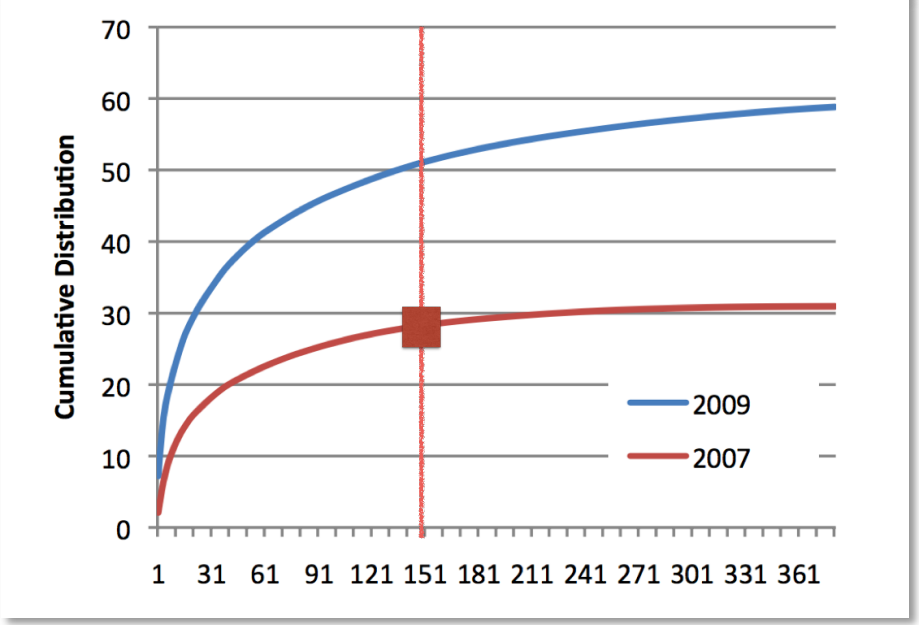

- In 2007, thousands of ASN make up 50% of traffic
- By 2009, 150 grouped ASN contribute 50% or more of traffic
- Data from 150 ISPs participating Internet Observatory (SIGCOMM 2010)

Internet Trends

2013

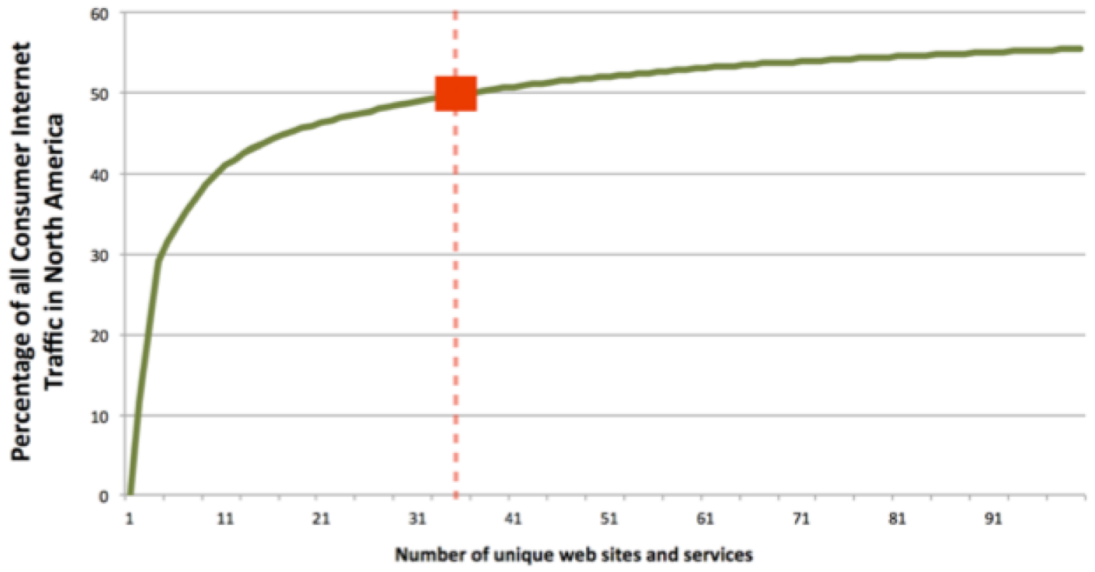

- On average, 35 ASN generate 50% of consumer traffic in 2013
- CDF of ingress peak traffic to participating providers in North America in May 2013

Internet Trends

2016

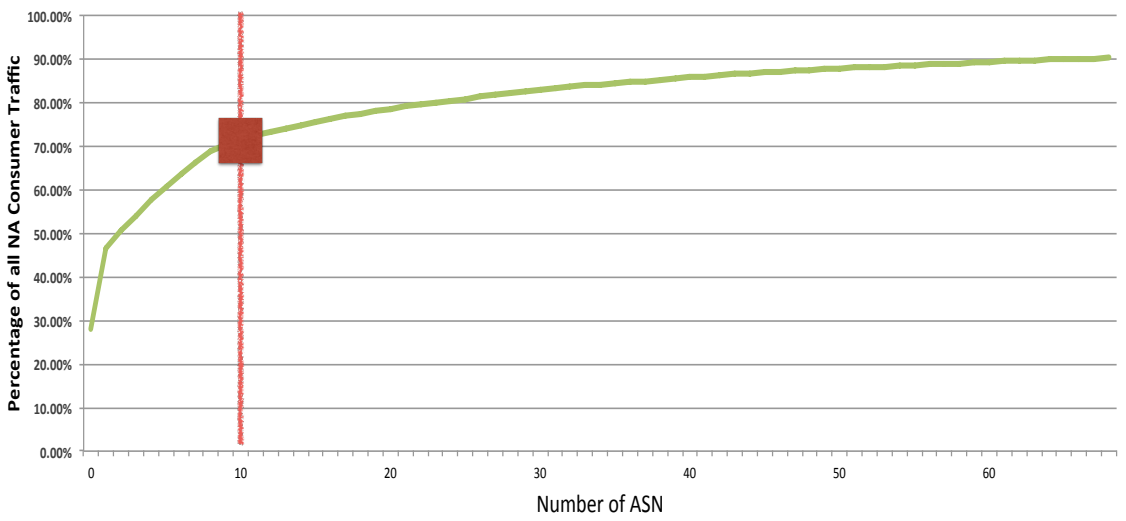

- On average, 10 ASN generate 70% of consumer traffic in 2016
- 30 ASN contribute more than 80% of all traffic
- Extremely heavy tailed distribution

Internet Trends - IP Counts

April 2016 (24h)

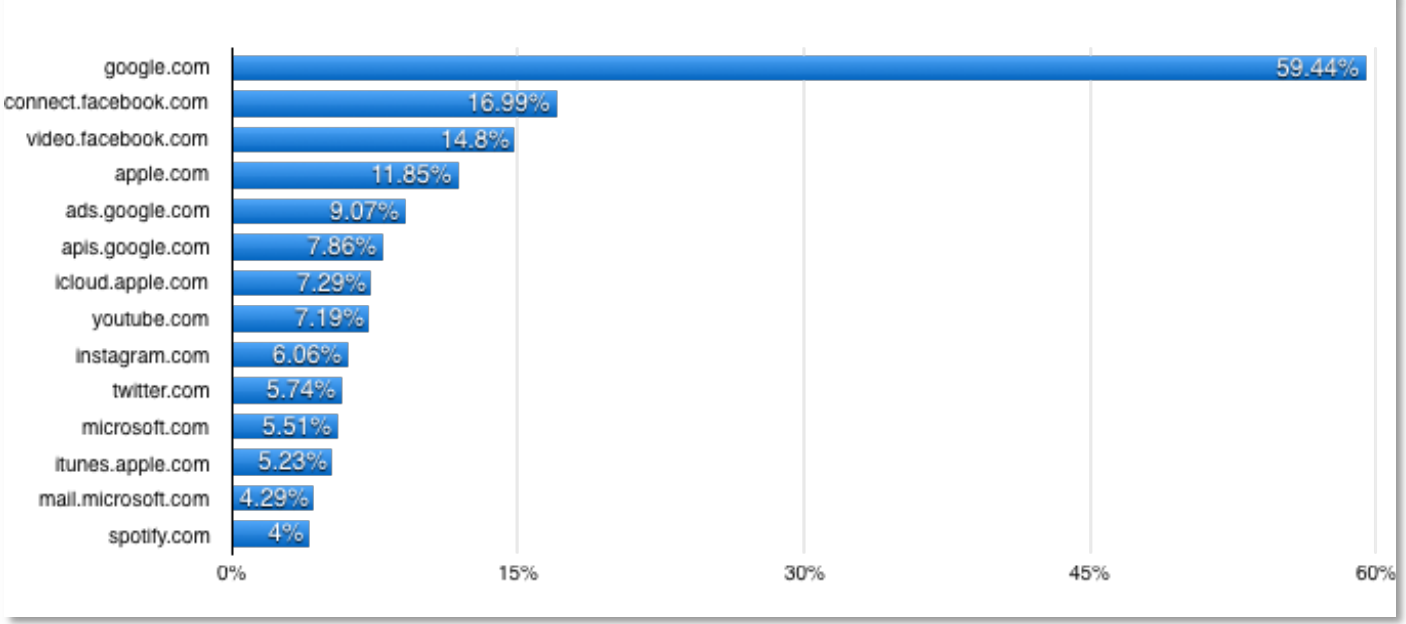

- Percentage of all consumer network IPs on average that send / receive traffic with service via direct or third-party (i.e. CDN / hosting) during 24 hour period during month of April 2016

Internet Trends – "The New Internet"

New Interconnections

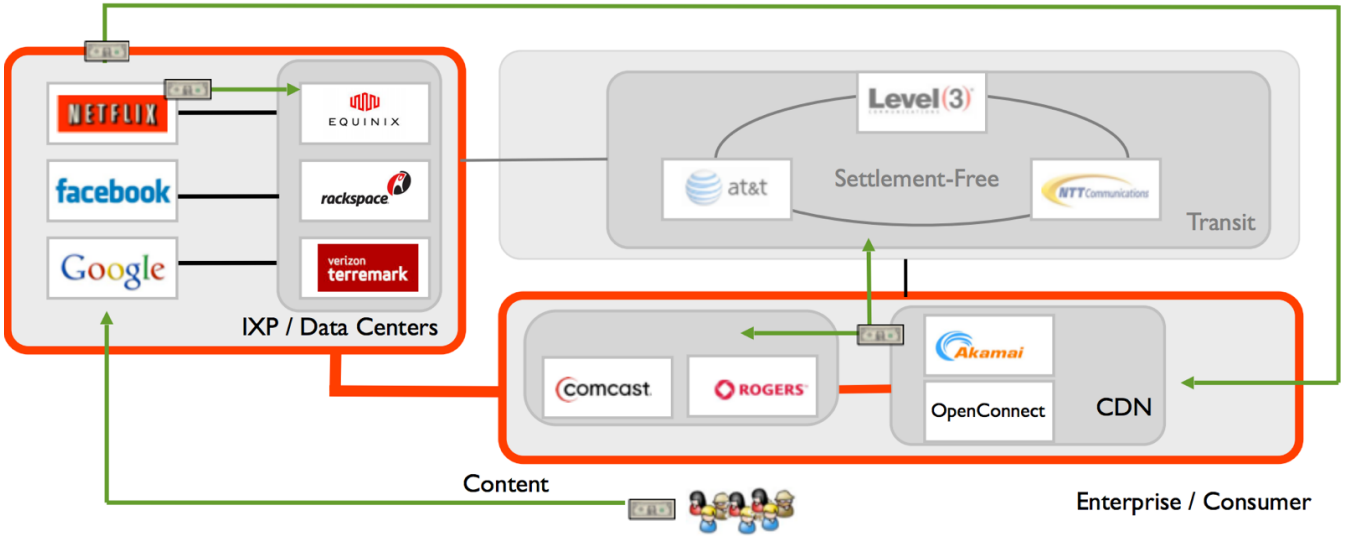

- Increasing volumes of peering at IXP and edge data centers. Significant growth in secondary markets. Growing consolidation in content delivery and sources.

Growth in CDN Traffic

CDN as a percentage of peak ingress traffic

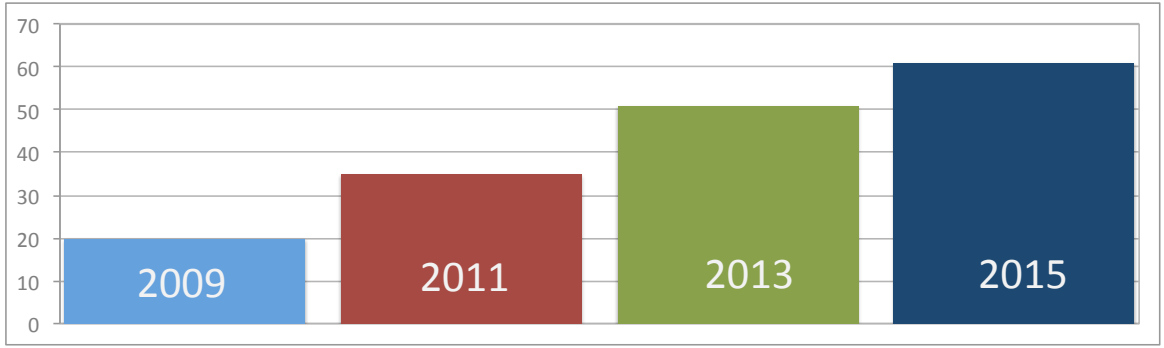

Data

- In 2009, CDN helped to offload less than 1/4 traffic. Most content delivered via peering / transit
- By 2015, the majority of traffic is CDN delivered from regional facility or provider based appliance
- (We define CDN as sole or multi-tenant distribution infrastructure deployed in multiple regional facilities or within multiple provider facilities. The definition is admittedly becoming increasingly fluid)

CDN Traffic 2016

CDN Traffic

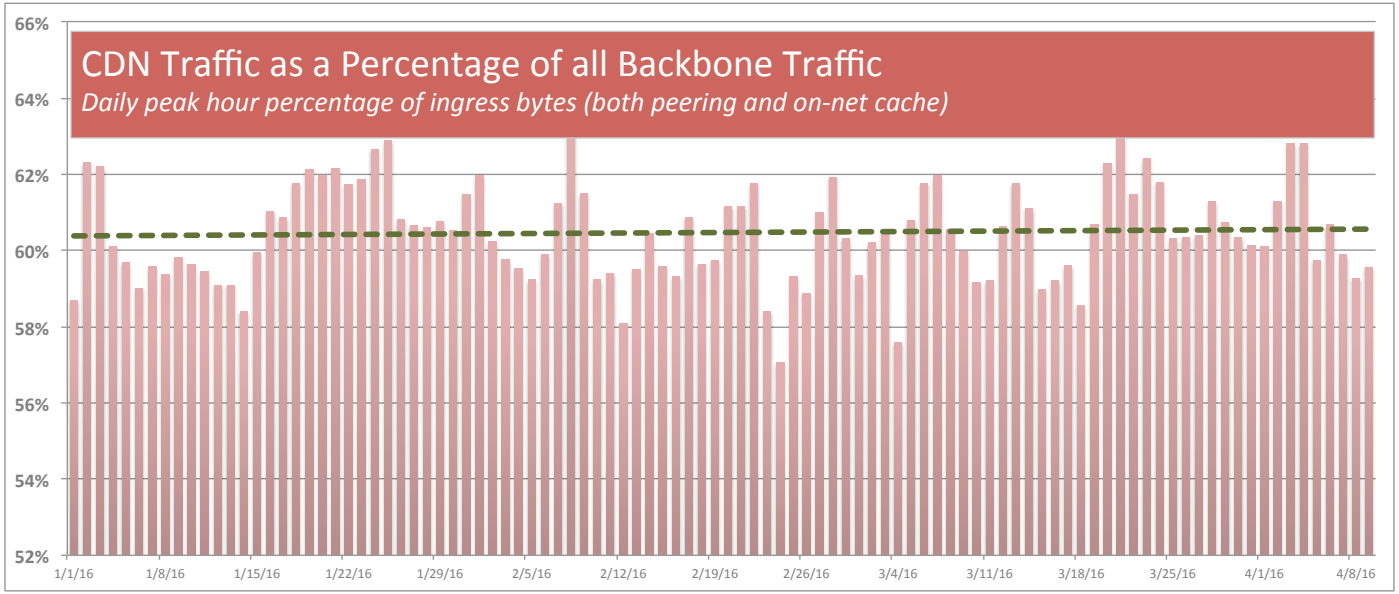

Data

- CDN growth as percentage of Internet traffic has slowed (1-2% a year)
- Likely due to market maturation (i.e. almost all video content is now CDN delivered)

CDN Traffic 2016

CDNs

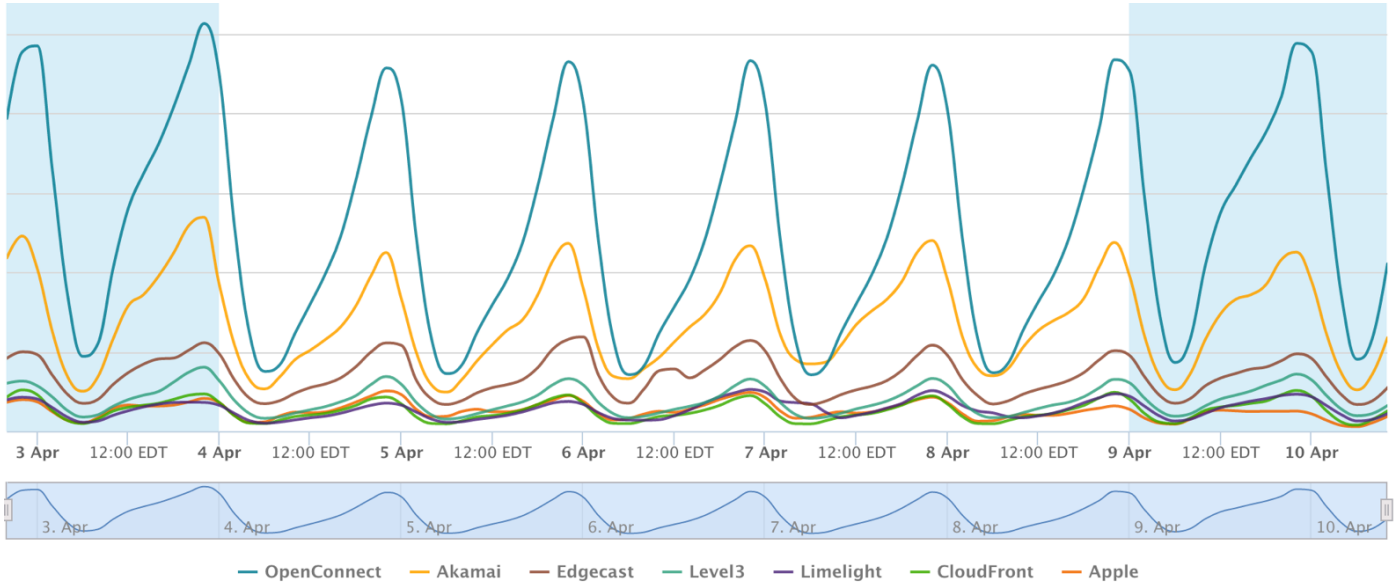

Data

- Anonymized, aggregate data from US providers
- OpenConnect is largest CDN at peak but GGC (where deployed) has larger daily aggregate volume for YouTube

CDN Traffic 2016

2016 CDN Observations

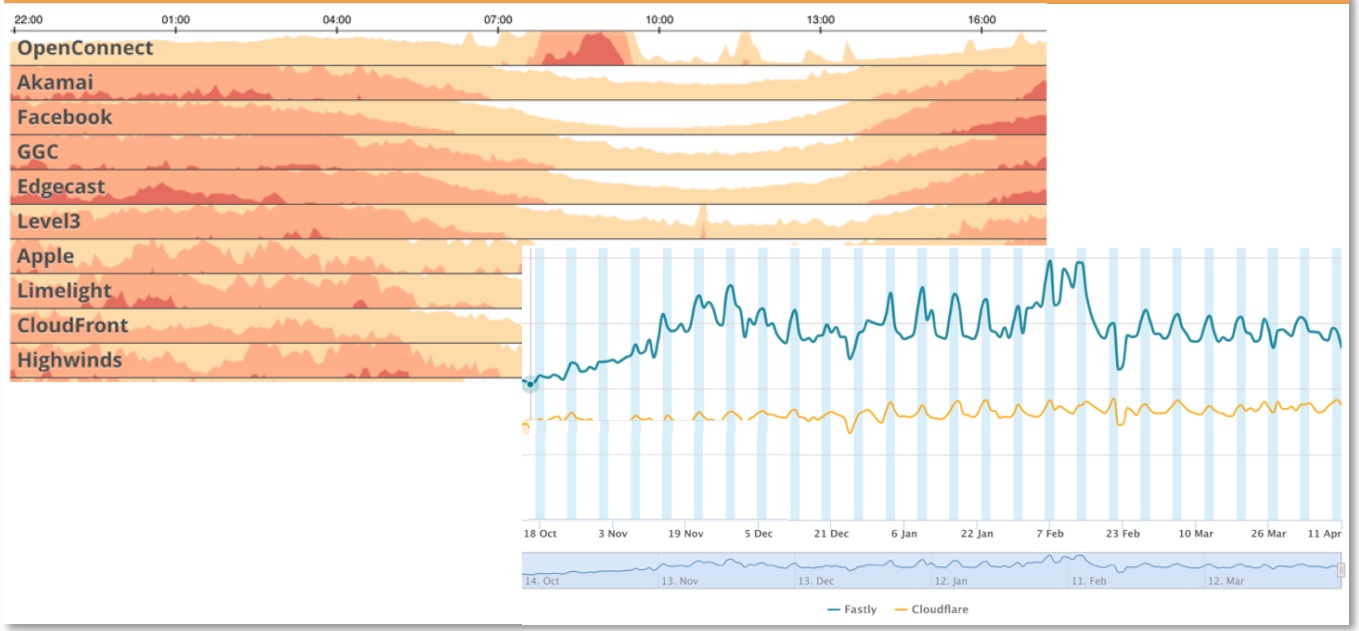

- CDN is the network
- Most traffic is adaptive bitrate
- Coordinating cache fill times becomes important and cache placement a growing issue
- Growing commercial agreements around CDN delivery and efficacy (e.g. locality)
- New CDN entrants (e.g. Fastly, CloudFlare, CloudFront) gaining market

Encrypted vs Un-Encrypted - Netflix Delivery Comparison

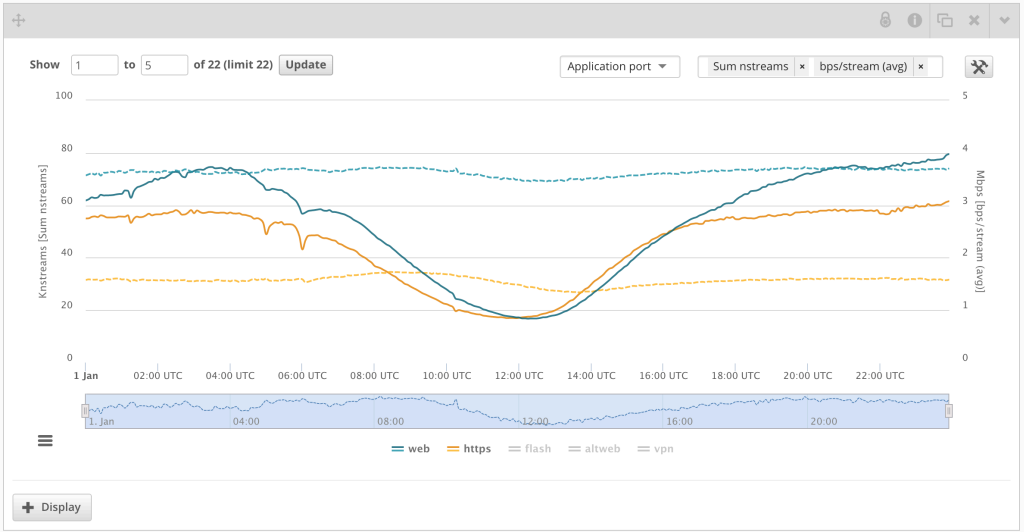

Compare performance and usage of encrypted vs. unencrypted Netflix traffic

Summary

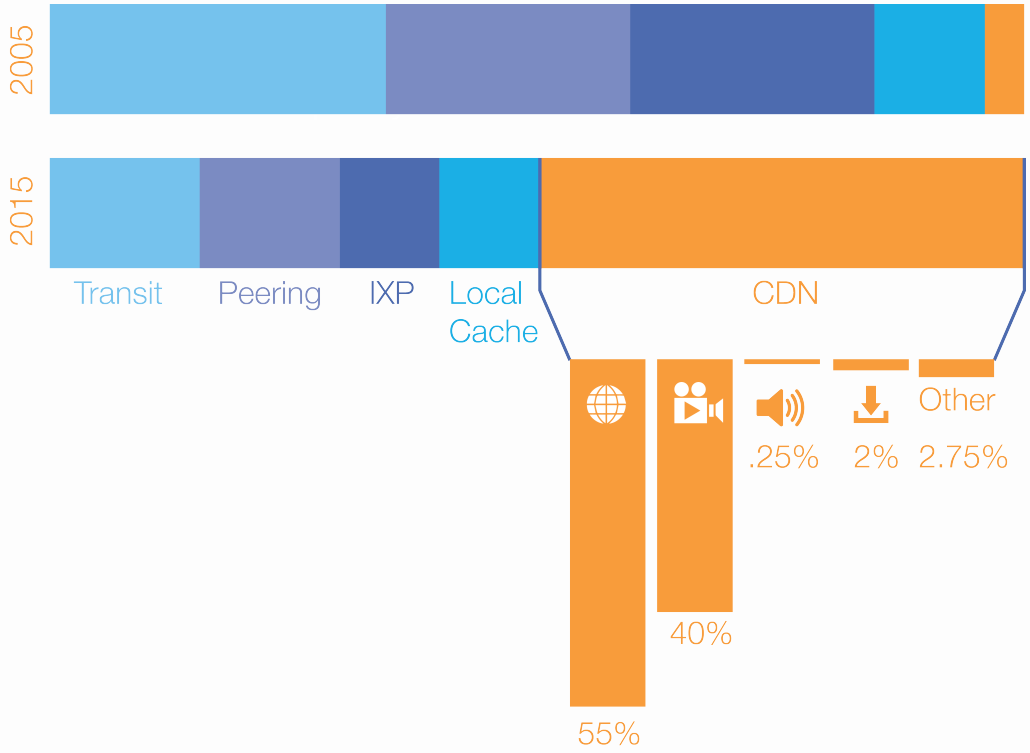

Internet Usage Evolution
and how content is delivered

- Today 30 sites generate 60% user traffic
- High Value Traffic is delivered from CDNs
- IXPs are Game-Publishers choice
- Transits deliver “Other”
- Adult Content

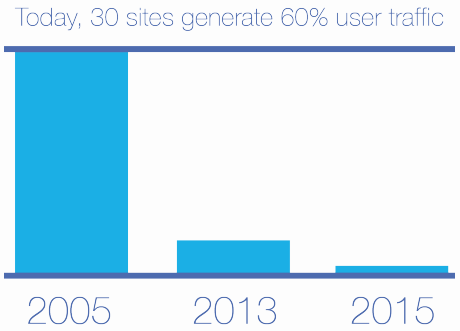

Deepfield Platform

Horizontally scalable, fully self-contained software architecture.
Download and run on your servers or use Deepfield SaaS.

Connectors	Network Logic	Service Map Logic	Real-Time Streaming
Query Language	Machine Learning	Alerting	Visualization

Deepfield Connectors – Any Data, Any Size

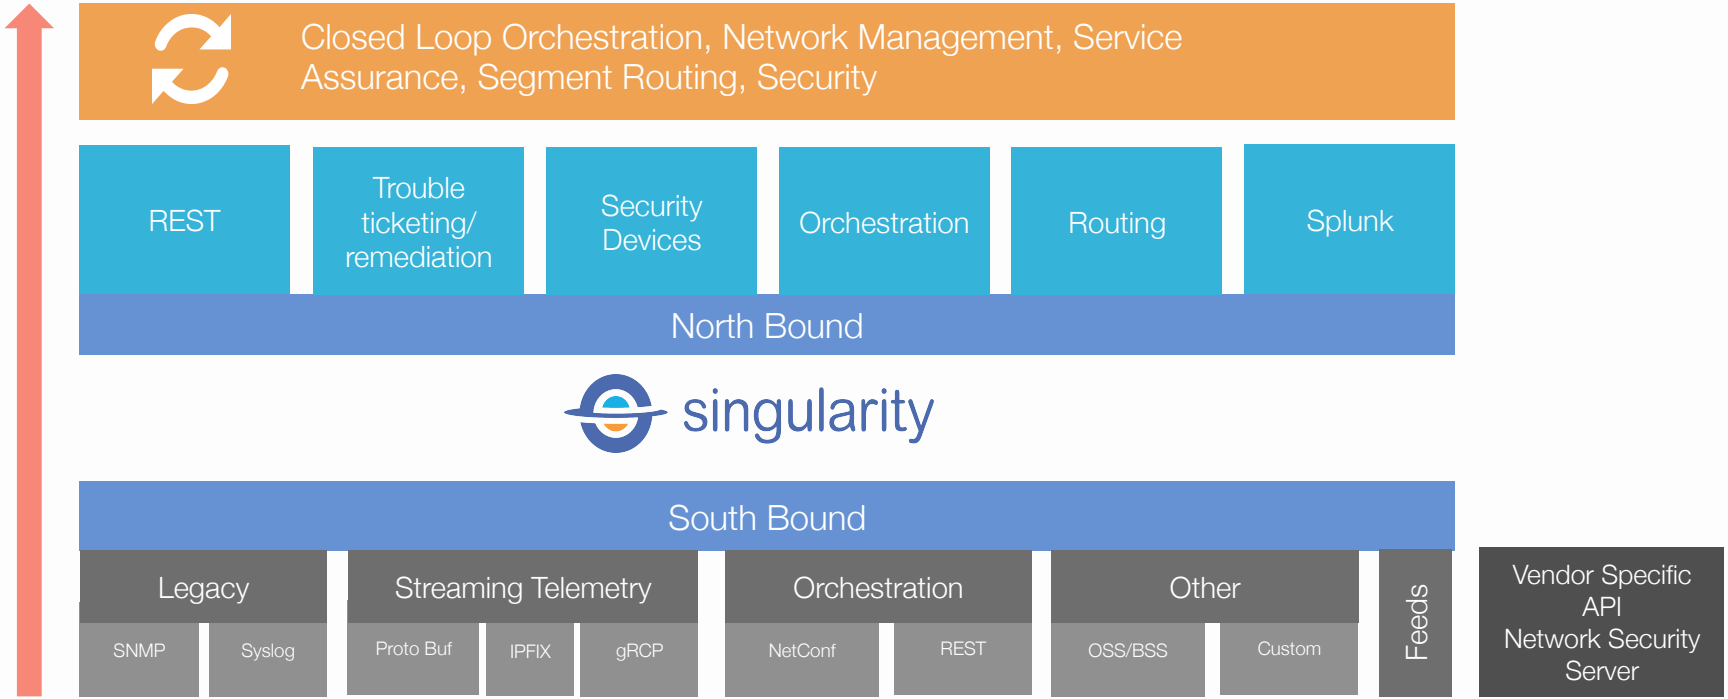

Wide-Angle Correlation

Deepfield logic engines correlate every session and subscriber to every network and data center component

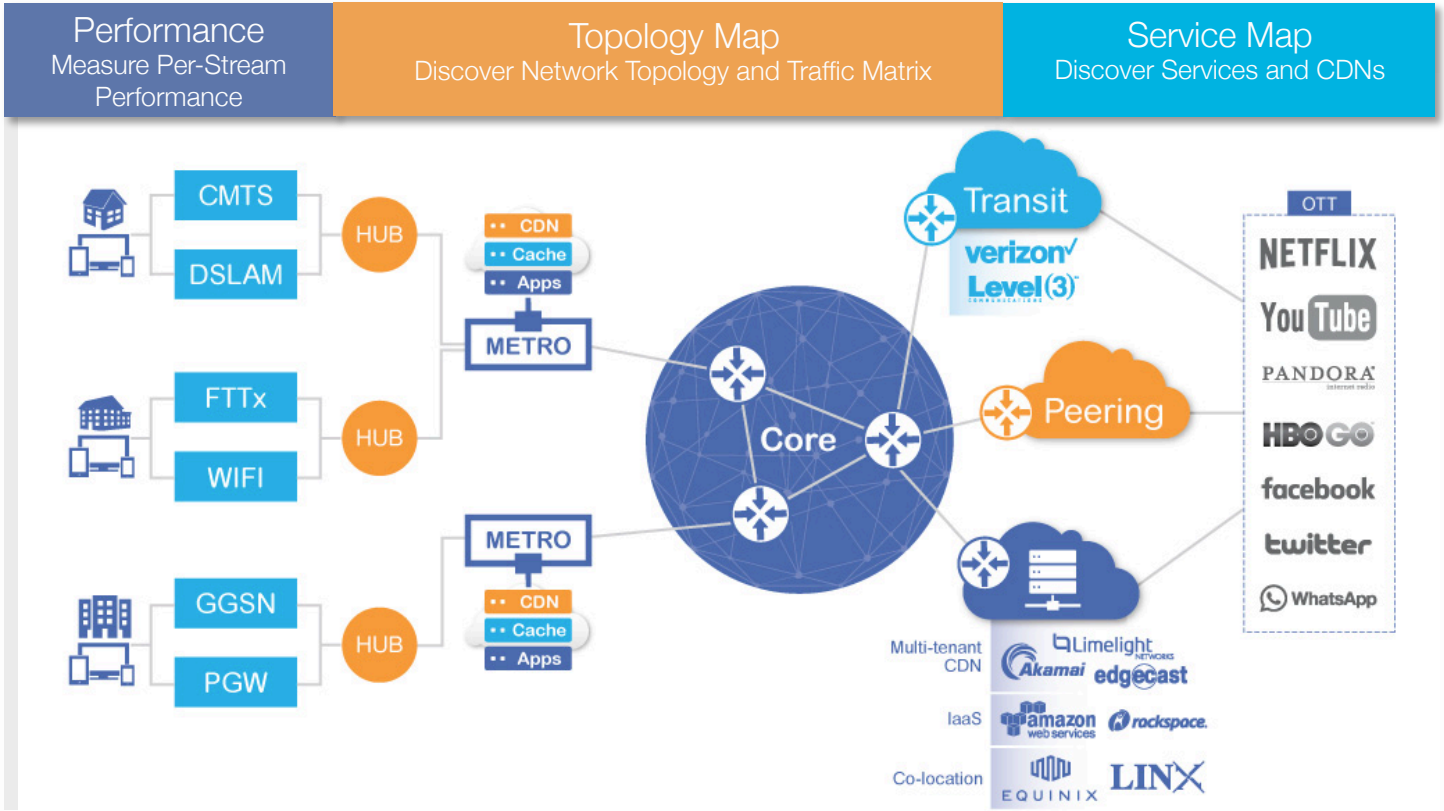

Questions?