# Modern Interconnects in the Internet Ecosystem

TELECOMS DIGITAL STRATEGY FORUM

Suva, Fiji

15th & 16th October 2018

Philip Smith

# A Bit of History…

- In the beginning, there was no Internet Backbone
  - Operators of the early networks just interconnected..

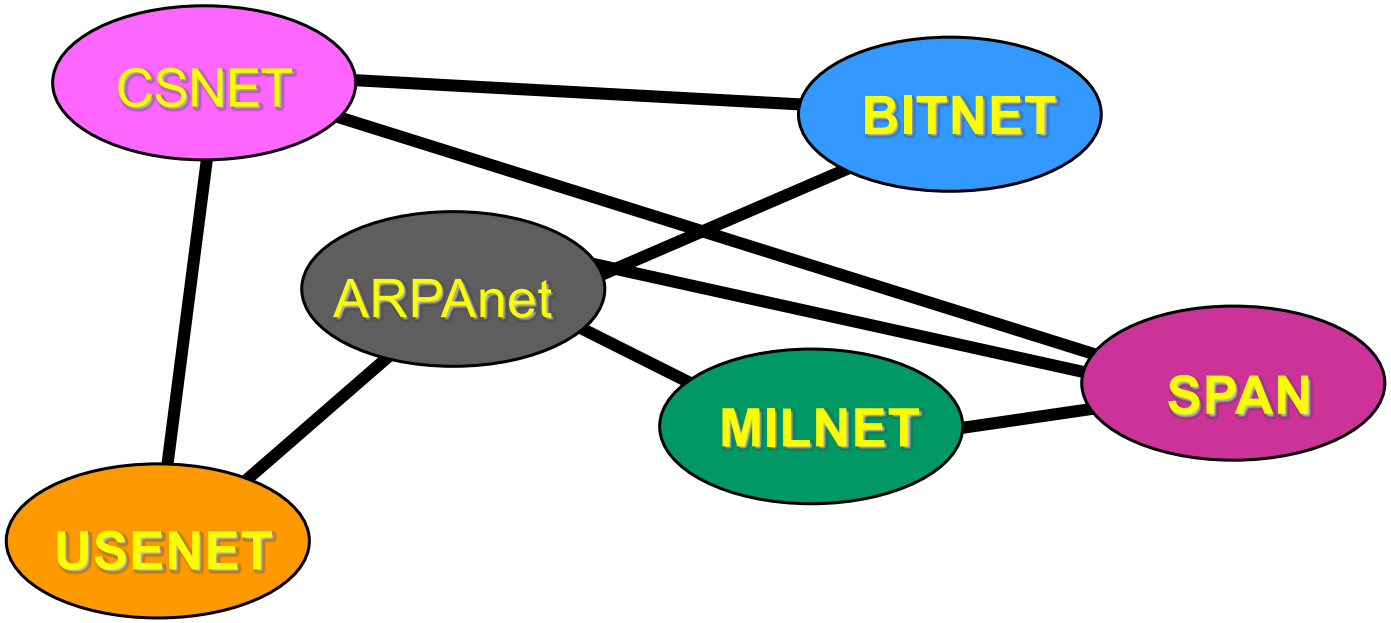

# The Old NSFNET Backbone

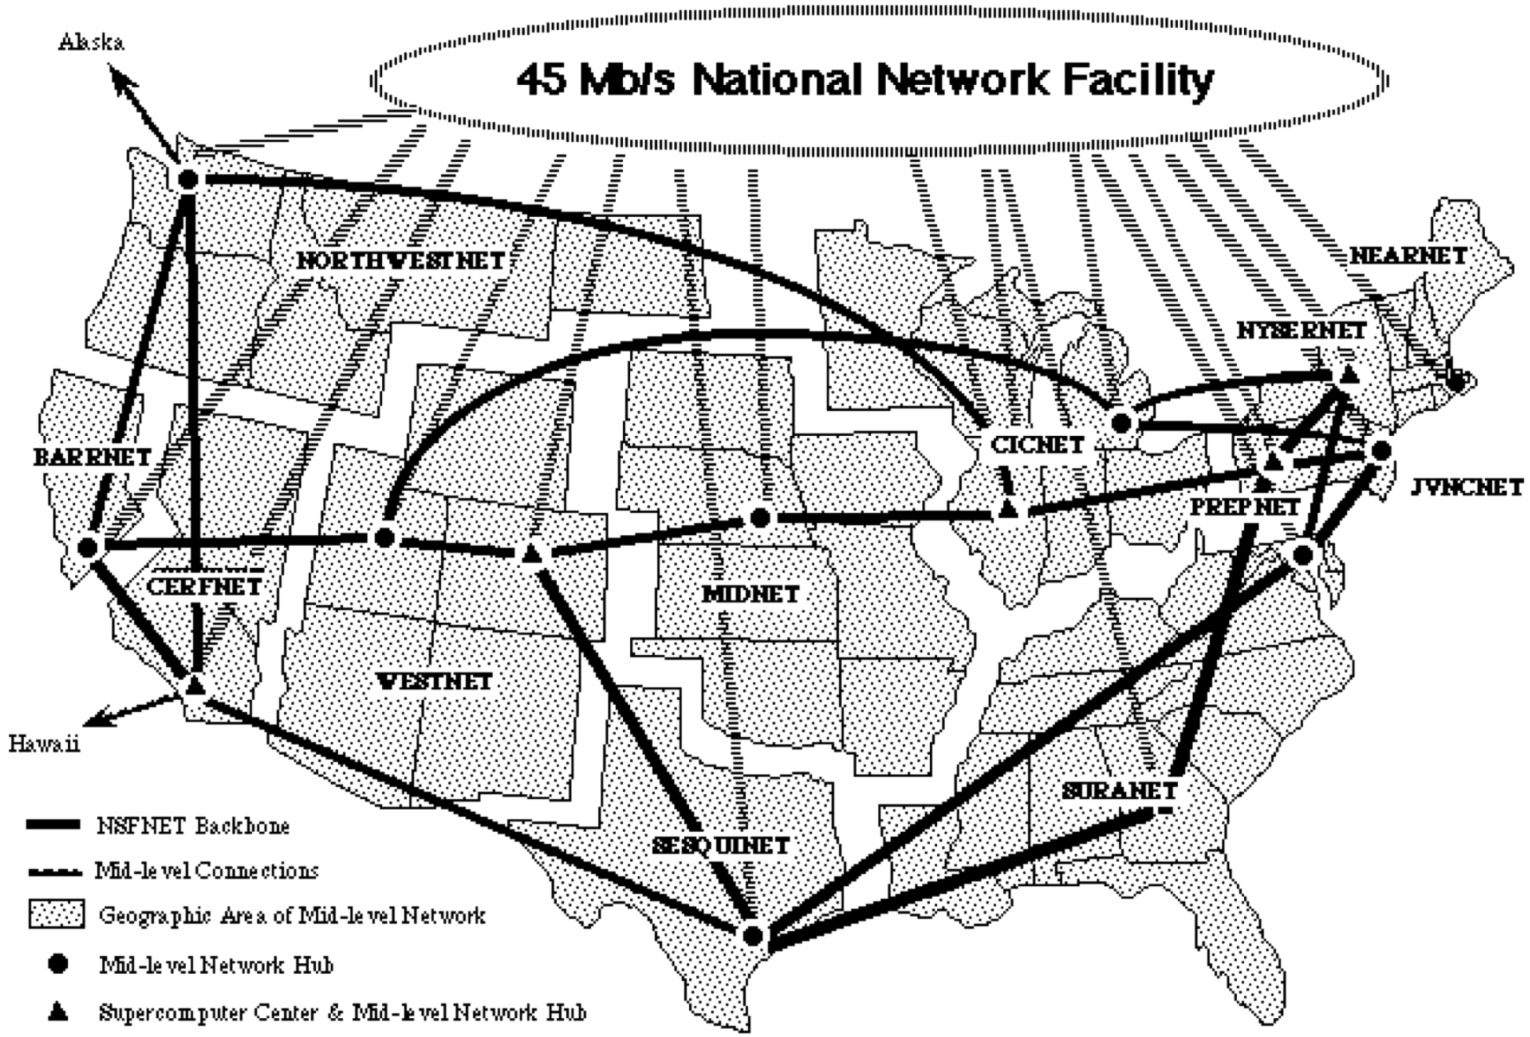

# Internet in the 1990s

- By mid-1990s, Internet model looked like this:
  - Very much US centric
  - NSPs provided transit coast-to-coast across the US
  - IXPs provided the interconnects in key centres

- NSPs of the mid-1990s became known as Tier-1s
  - Tier-1 is a network operator who has no need to buy transit from any other operator
  - Interconnect with other Tier-1s by Private Interconnect

# Tier-1 Private Interconnects

- "ANS, MCI and Sprint Sign Agreements for Direct Exchange of Internet Traffic" – June 30, 1995

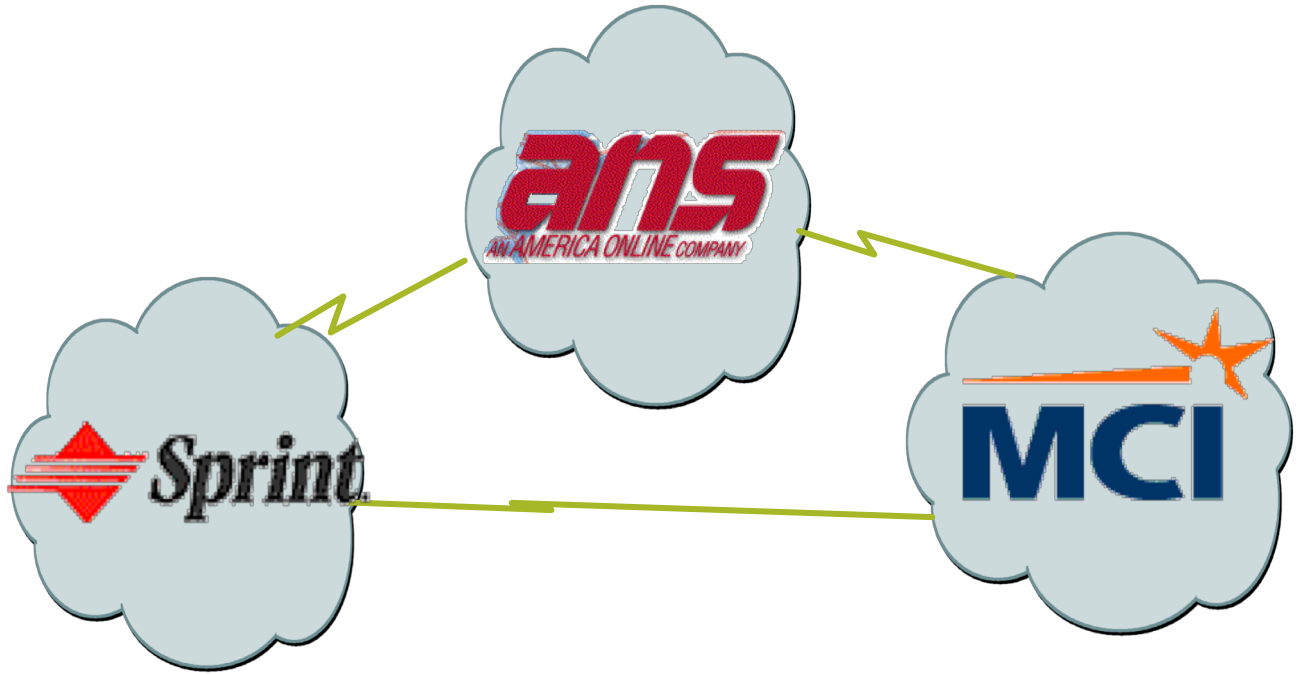

# The Internet Today

- “Content is King”
- The typical end-user traffic profile shows:
  - 50% of all Internet traffic is Google/YouTube
  - 25% of all Internet traffic is Facebook
  - 10% of all Internet traffic is Content hosted by Akamai, Cloudflare, Netflix, Microsoft, and other content operators
  - (“typical” in this author’s experience)

- This is a significant change over the traffic profile from the late 1990s and early 2000s

# The Internet Today

- Major content distribution networks no longer have “one big server”
- They each operate a substantial distributed network of content delivery caches from multiple regional datacentres
- Goal:
  - Content as close to the “eyeballs” (the end users) as possible
  - Lowest latency possible
  - Highest bandwidth possible
- The average consumer’s tolerance of non-working websites or delays is only a few seconds

# Global Internet: High Level View

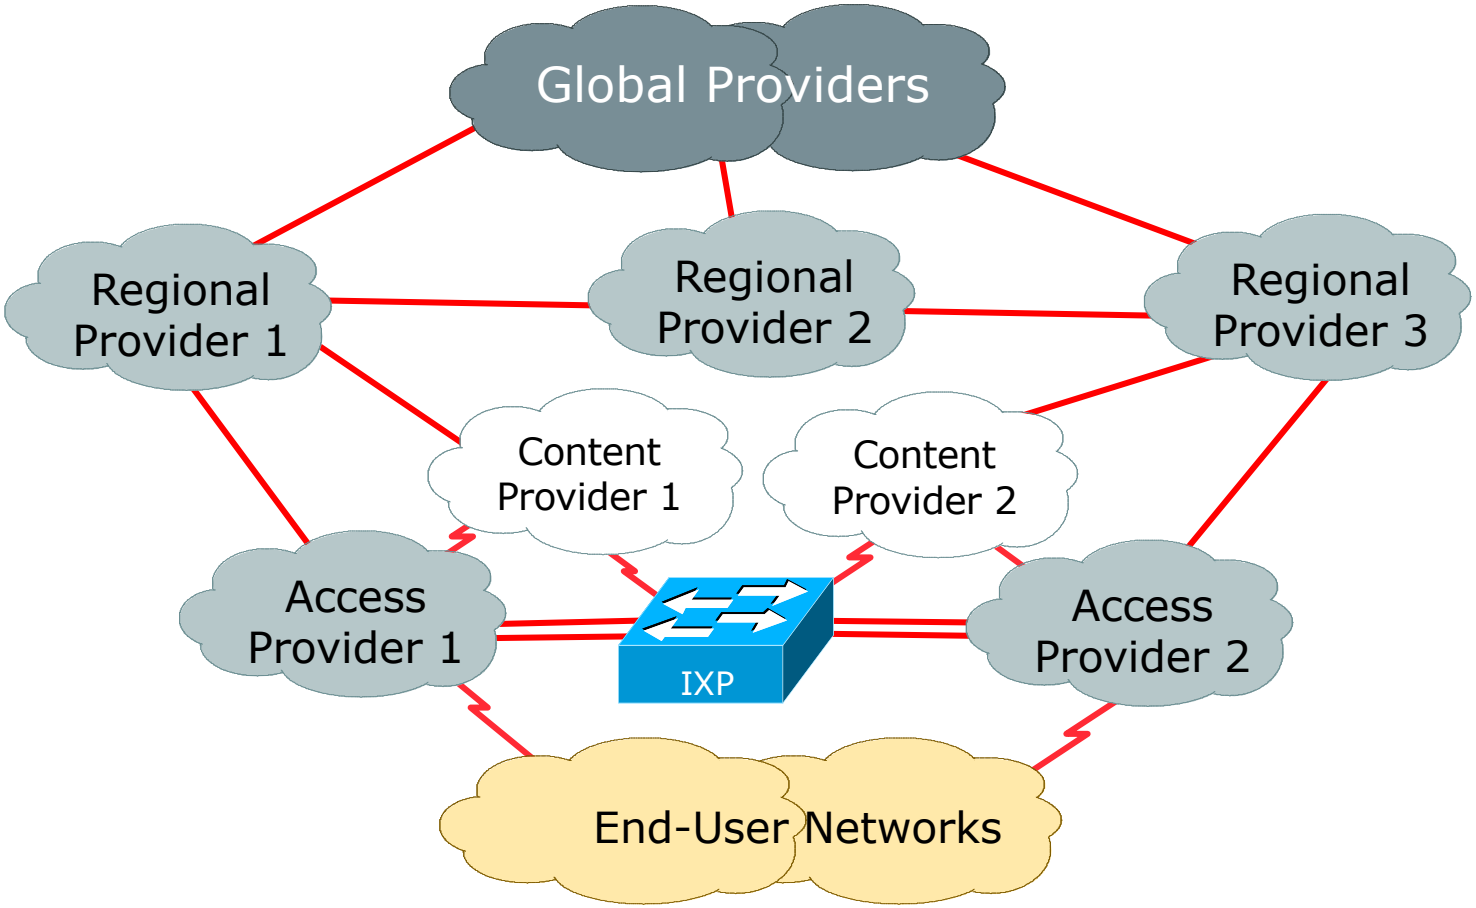

# Internet Provider Profile

- Content Providers have moved close to the Access Providers and to Public Interconnects
- Access Providers are simply a vehicle to deliver content as fast as possible to end-user
- Content Providers directly connect with Access Providers
  - PNI – Private Network Interconnect, *or*
  - Across Internet Exchange Points (IXPs), *and*
  - Provide a local cache for most frequently used content

# Content delivery is competitive!

- Competition in local marketplace is all about speed and quality of content delivery
  - e.g.

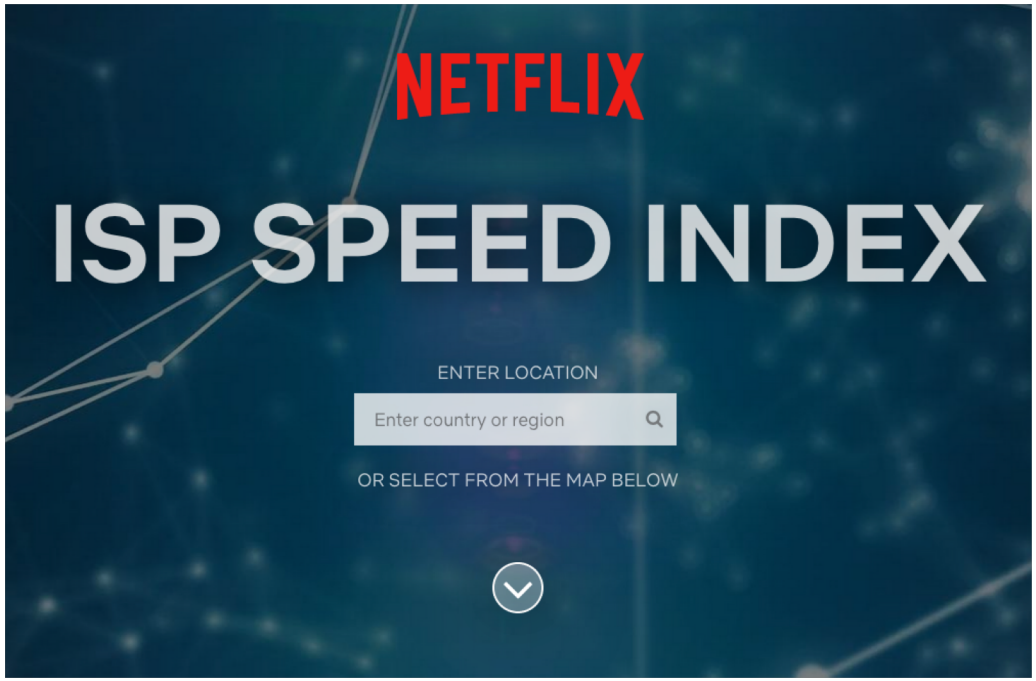

# What happened?

- In the late 1990s:
  - US was hub of global Internet
  - Europe was becoming a hub of the European Internet
  - Asia, Pacific, Latin America still mostly connected to the US, rather than interconnected within region
  - Africa mostly connected to Europe, rather than interconnected within region

  - Internet access was by desktop or, more rarely, laptop computer
    - Content by static web pages, UseNet, some news media
  - No smartphones or tablets or 3G or LTE

# What happened?

- Apple iPhone launch in January 2007
  - Availability of 3G networks
  - Smartphones took off
  - Google's Android quick to follow
- Dominance of Google as search engine
- Dominance of Facebook for social networking

- By 2010, users could be online 24x7 through their increasingly smarter and more data-hungry devices

# Content Distribution Today

- CDNs such as Google, Facebook, Cloudflare and Akamai have built considerable content distribution infrastructure
- Several have large stake holdings in global submarine fibre
  - Example: https://www.wired.co.uk/article/google-facebook-plcn-internet-cable
- Several have built their own large data centres at strategic locations around the globe
- Replaced the Tier-1 operator as the content delivery vehicle to the regions around the globe
- CDNs encourage operators to connect to their datacentres to maximise performance for content delivery

# Content Distribution Today

- CDNs such as Google, Facebook and Akamai also supply and operate content caches
- Operators with a few Gbps of content being served from these CDNs usually qualify for a cache
- Caches are found in most larger operators today

- Many IXPs have CDNs present
- Many operators at smaller IXPs will share their content caches with their peers across the fabric

# Content Distribution Today

- CDNs at IXPs:
  - Lowest possible latency between the content and the end-user
  - Highest possible bandwidth between the content and the end-user

  - A Happy End-user! End-user stays connected to the CDN operator, rather than moving to a competitor

  - Onus on network operator to maintain high capacity at IXP and on to end-user
    - International connectivity is usually much more expensive!

# CDN

- Example of CDN Cache effectiveness
- Feeding over 3 times more than is arriving via transit link
- Peers benefiting

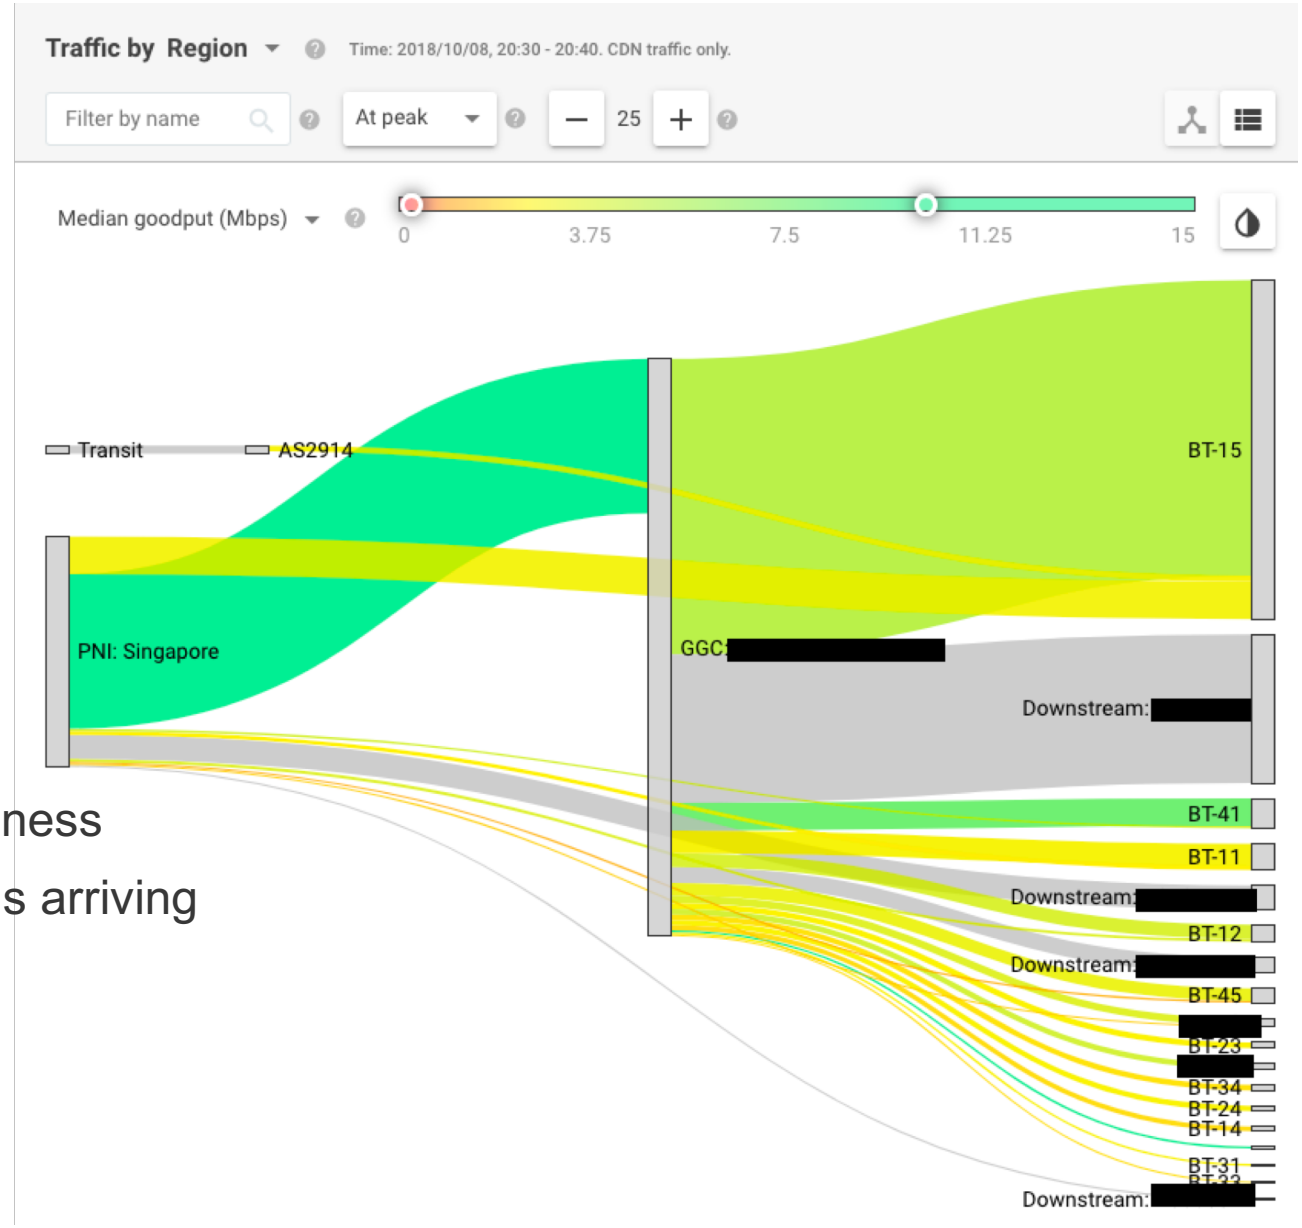

# Network Operator Goals?

- Today, the vast majority of content consumed by end-users is available by peering:
  - The major content providers (Google, Facebook, etc)
  - Private cross connects
  - Internet Exchange Points
- A network operator's goal is to obtain as much peering as possible
- Transit is for the last resort, for any content not available by peering

# Network Operator Goals?

- Peering
  - Locally with direct cross-connect with other providers
  - Locally at an Internet Exchange Point
  - Getting to the topologically nearest IXP or other interconnect

- Transit
  - Relying on another network operator to get the rest of the Internet
  - Considered a last resort now

# New Technologies

- Network Operators have designed their networks now to ensure reliable & high bandwidth delivery of content
  - From Caches
  - From CDN operators
  - This is the majority of their traffic

- These updated infrastructures facilitate:
  - New services (eg “Cloud Computing”)
  - 5G cellular infrastructure
  - Many more connected devices (eg “Internet of Things”)

# New Behaviours

- Peering and Interconnects are more important than they ever have been
  - IXPs and Private Interconnects already very popular in Europe and North America, and have been for 20+ years
  - Emergence of major interconnects in Asia and Latin America in the last decade

- Peering “meet-me” events:
  - Are a part of almost every single Network Operator Conference
  - Standalone events happen on Global, Regional & Local levels

# New Behaviours

- IXP establishment is critical for any domestic Internet economy
  - An open neutral location where any and every network operator can freely interconnect their networks to exchange traffic

- No IXP means
  - Costly interconnects
  - Congested infrastructure
  - High latencies
  - Poor quality of internet service
  - A lagging Internet economy
  - Inability to support modern network technologies & services

# Internet Exchange Point

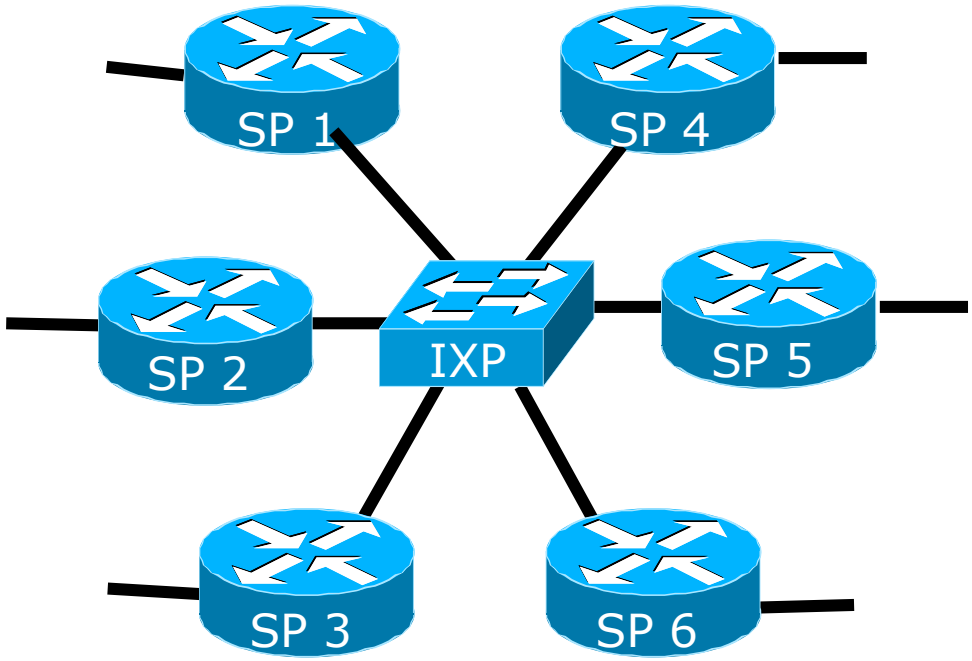


Each of these represents a border router of a different network operator

# IPv6

- New networks are deployed supporting dual stack
  - The infrastructure runs IPv6 and legacy IPv4 side by side
    - No interaction between IPv4 and IPv6 – independent protocols
- IPv4 address space is almost no longer available
  - Many network operators are now using private IPv4 address space (RFC1918 or RFC6598) and using Network Address Translation (NAT) to translate to public IPv4 addresses
- IPv6 address space is plentiful
  - IPv6 is supported on almost every networking device available today

# IPv4 & IPv6 dual stack operation

- IPv6 is designed to work independently of IPv4
- If a destination is available only over IPv4, IPv4 will be used
- If a destination is available over IPv4 & IPv6, Happy Eyeballs (RFC8305) ensures that the client uses the transport for the best user experience

# *Country experiences*

Lao PDR

Cambodia

Brunei Darussalam

Mongolia

Bhutan

**Annual (regional / sub-regional) training on IPv6 deployment and IPv6 Infrastructure Security 2011 onwards**

**Specialized technical advice and training to countries and interested telecom operators**

**Training during PACNOG 2017**

**Recommendations on IPv6 deployment**

COMMUNICATIONS REGULATORY COMMISSION OF MONGOLIA

INFORMATION TECHNOLOGY, POST AND TELECOMMUNICATIONS AUTHORITY

Australian Government

Department of Communications and the Arts

# Why not NAT?

- How to scale NAT performance for large networks?
  - Limiting tcp/udp ports per user harms user experience
- How to scale NAT throughput for large networks and new technologies such as LTE, LTE-A and 5G?
- NAT deployment requires redesign of SP network
- Network has to keep state of connections
- Breaks end-to-end network security
- Breaks non-NAT friendly applications
- Address sharing has reputation, reliability and security issues for end-users
- Makes the NAT device a target for miscreants due to possible impact on large numbers of users

# Implications for the Pacific

- Sydney and Los Angeles are the interconnect hubs for the Pacific
  - There are more optimum locations which offer much better RTT and performance than hauling traffic to/from/via Sydney and/or Los Angeles
- Example: The PacPeer Project explores optimum interconnections for network operators across the Pacific
  - https://pacpeer.org/
  - https://pacpeer.org/presentations/brewerj_peering_strategy_pacific_pacnog18.pdf

# Implications for the Pacific

- Fiji could be the regional hub for the South Pacific
- Guam could be the regional hub for the North Pacific
- Both Fiji & Guam have:
  - Large amounts of submarine fibre passing through
  - No open neutral interconnect facility
- Hawaii should be the regional hub for the whole Pacific
  - (following the fibre paths)
  - But capacity is cheaper direct to Los Angeles (even though latency more than doubles)
    - (Pacific to Hawaii + Hawaii to Los Angeles is more expensive than Pacific to Los Angeles)

# Pacific Fibre

Submarine fibre map

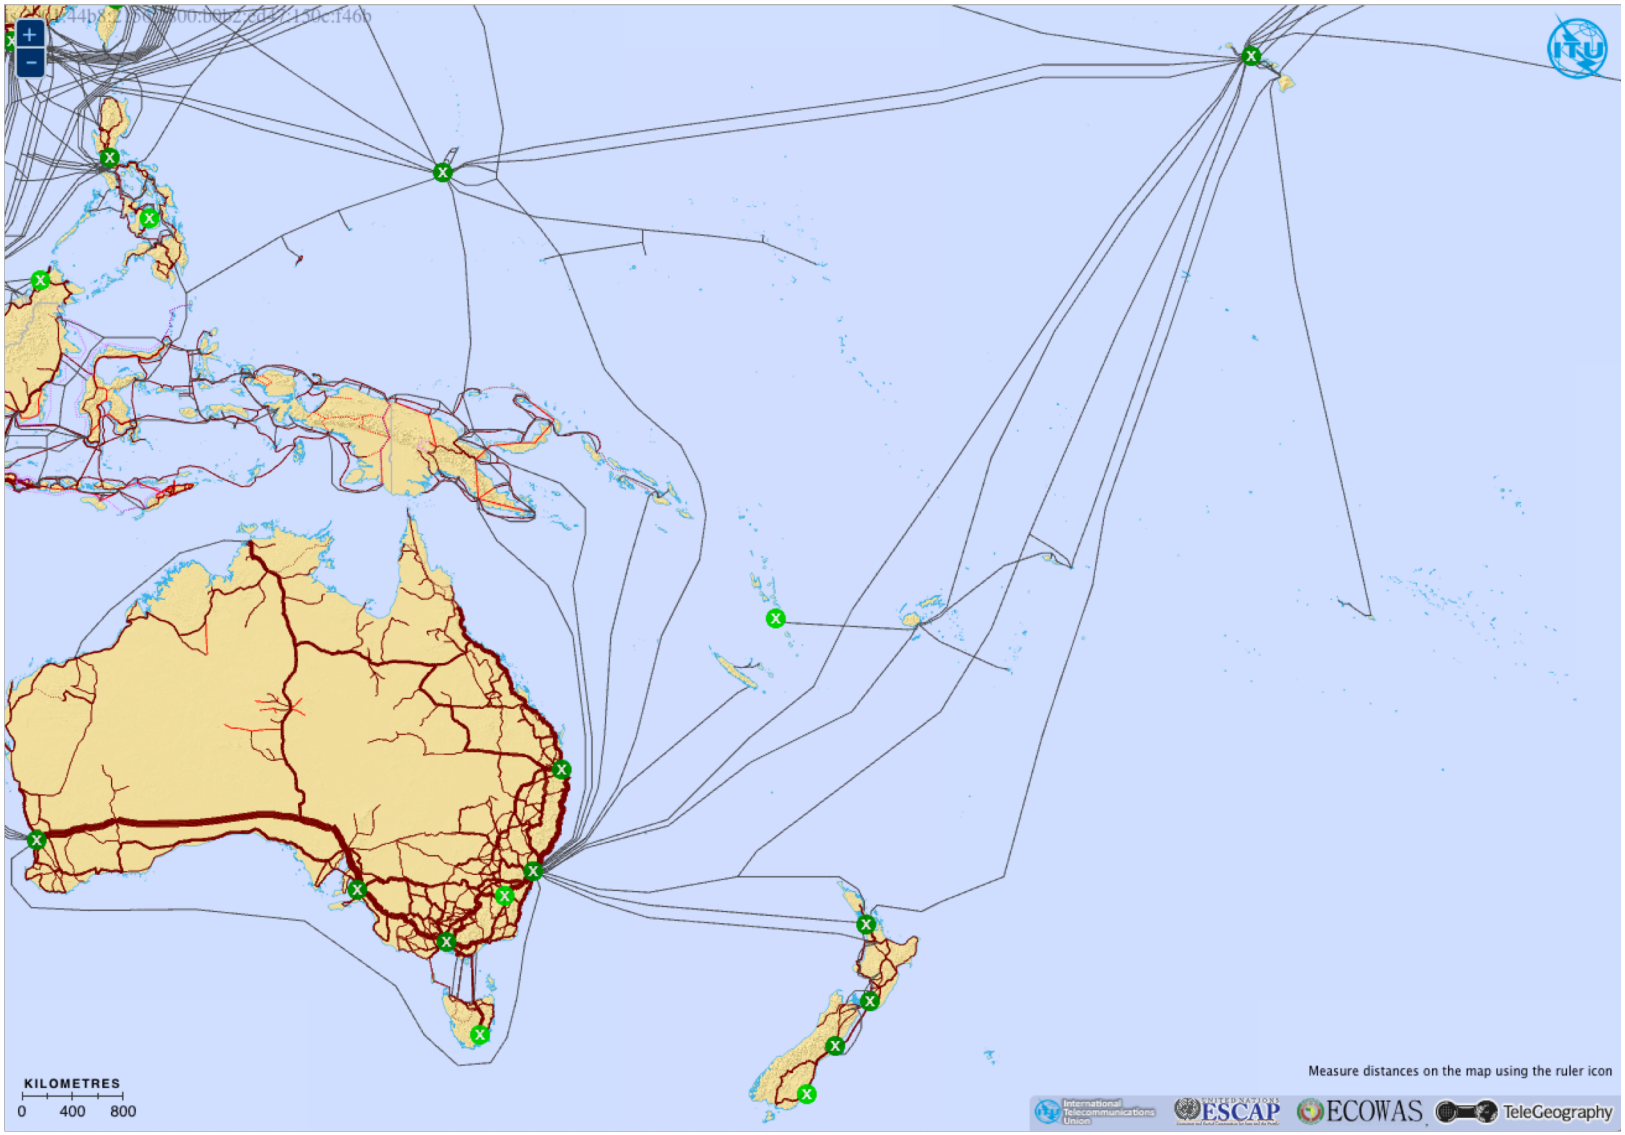




**Source: https://www.itu.int/itu-d/tnd-map-public/**

# Implications for the Pacific

- Apart from Fiji, Guam, and Hawai'i becoming Regional Hubs…
- What are the other priorities?
  - The network operators operating in a nation need to interconnect their networks
    - Inexpensive, high bandwidth, low latency = best end-user experience
  - Every single nation which has more than two network operators needs to have an open neutral interconnect
    - Not just a closed interconnect for a select few
- Success story: Vanuatu established the first IXP in the Pacific (outside Australia, New Zealand and Hawai'i)

# Evolution Summary: 20 years ago

- Centralised Internet (in US & Europe)
- Very diverse content, and hosted at origin
- Clear hierarchy of Tier-1s, Regional providers, and Access providers
- No 3G data networks or smartphones

- Access provider goal was to provide international connectivity to content and email

# Evolution Summary: Today

- Model of centralised Internet is no more
- “Content is King”
  - >80% of traffic volume is from the major content providers
  - Network operator focus today is on delivering content from the major content providers more efficiently than their competitors
  - CDN “performance meters” and “SpeedTests” are now customer measures of Internet Quality of Service
  - Latency, bandwidth & optimisation of regional traffic is top priority
- Geoff Huston opinion piece:
  - https://blog.apnic.net/2016/10/28/the-death-of-transit/

# Thank You