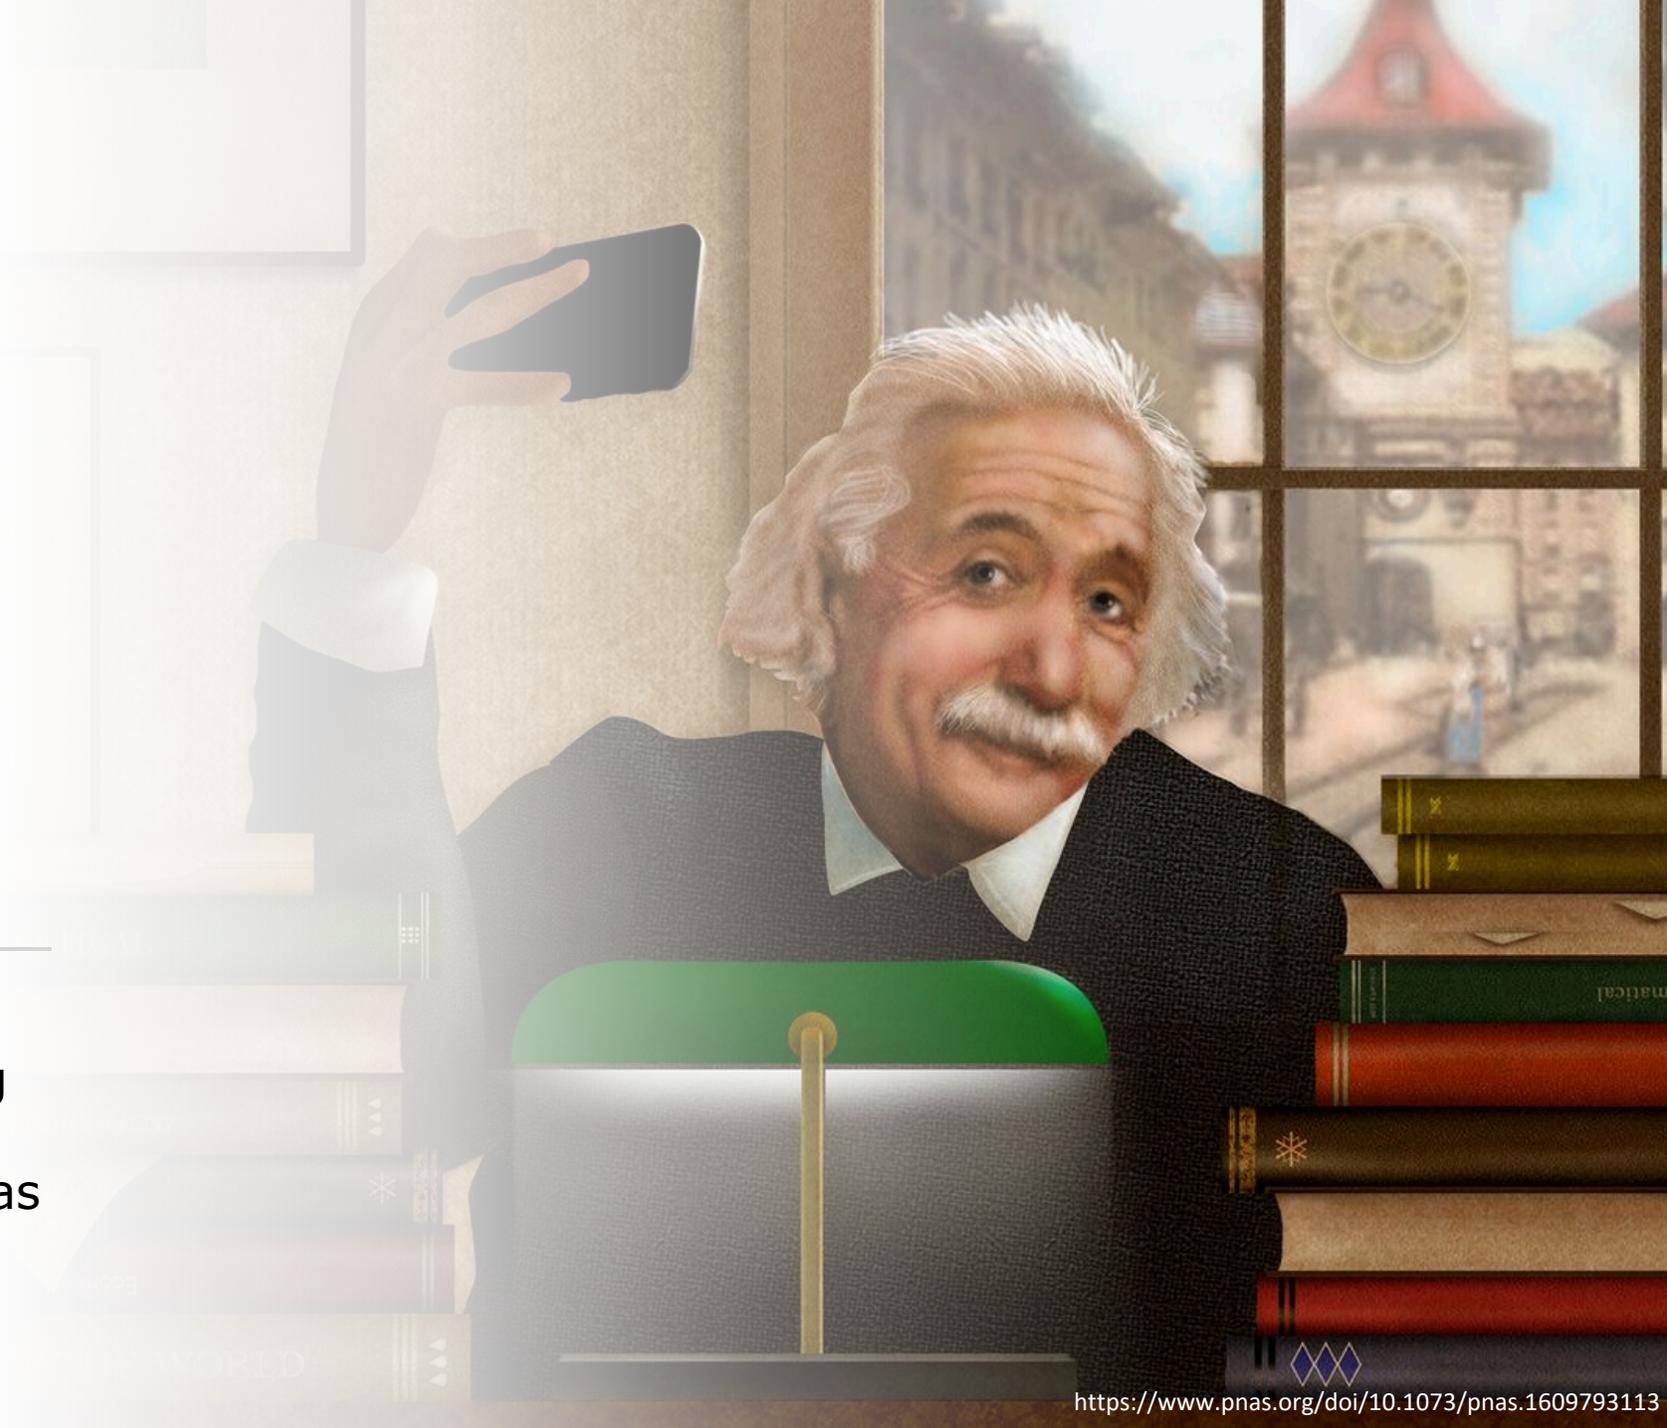


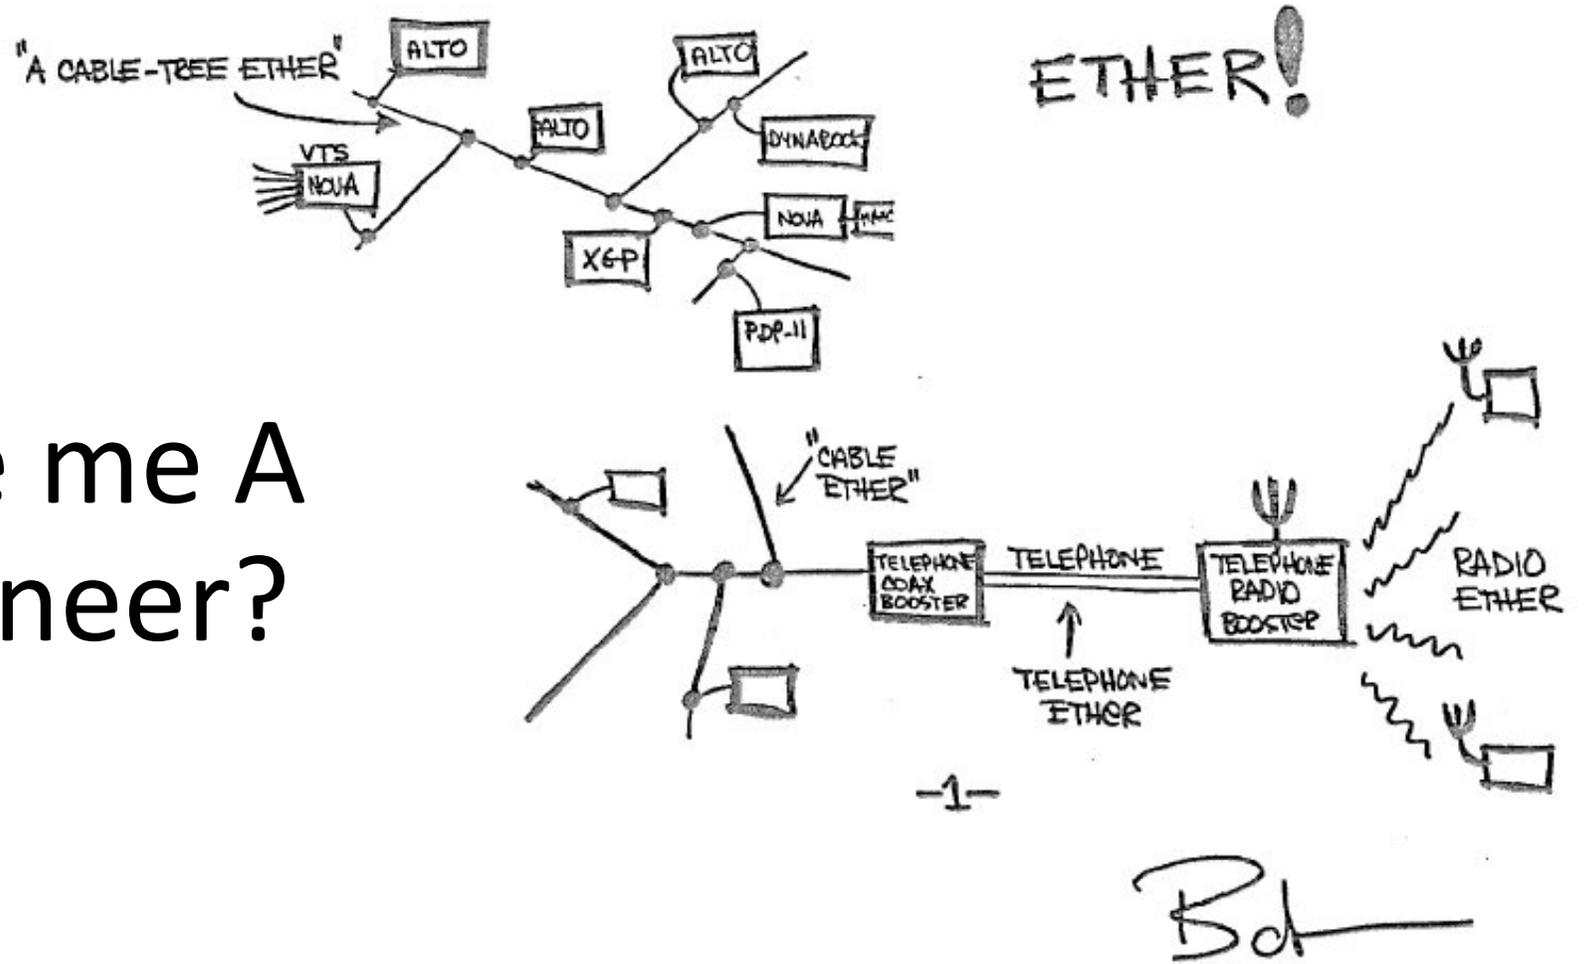
—

“The internet
is fine – I’m
on Facebook
right now!”

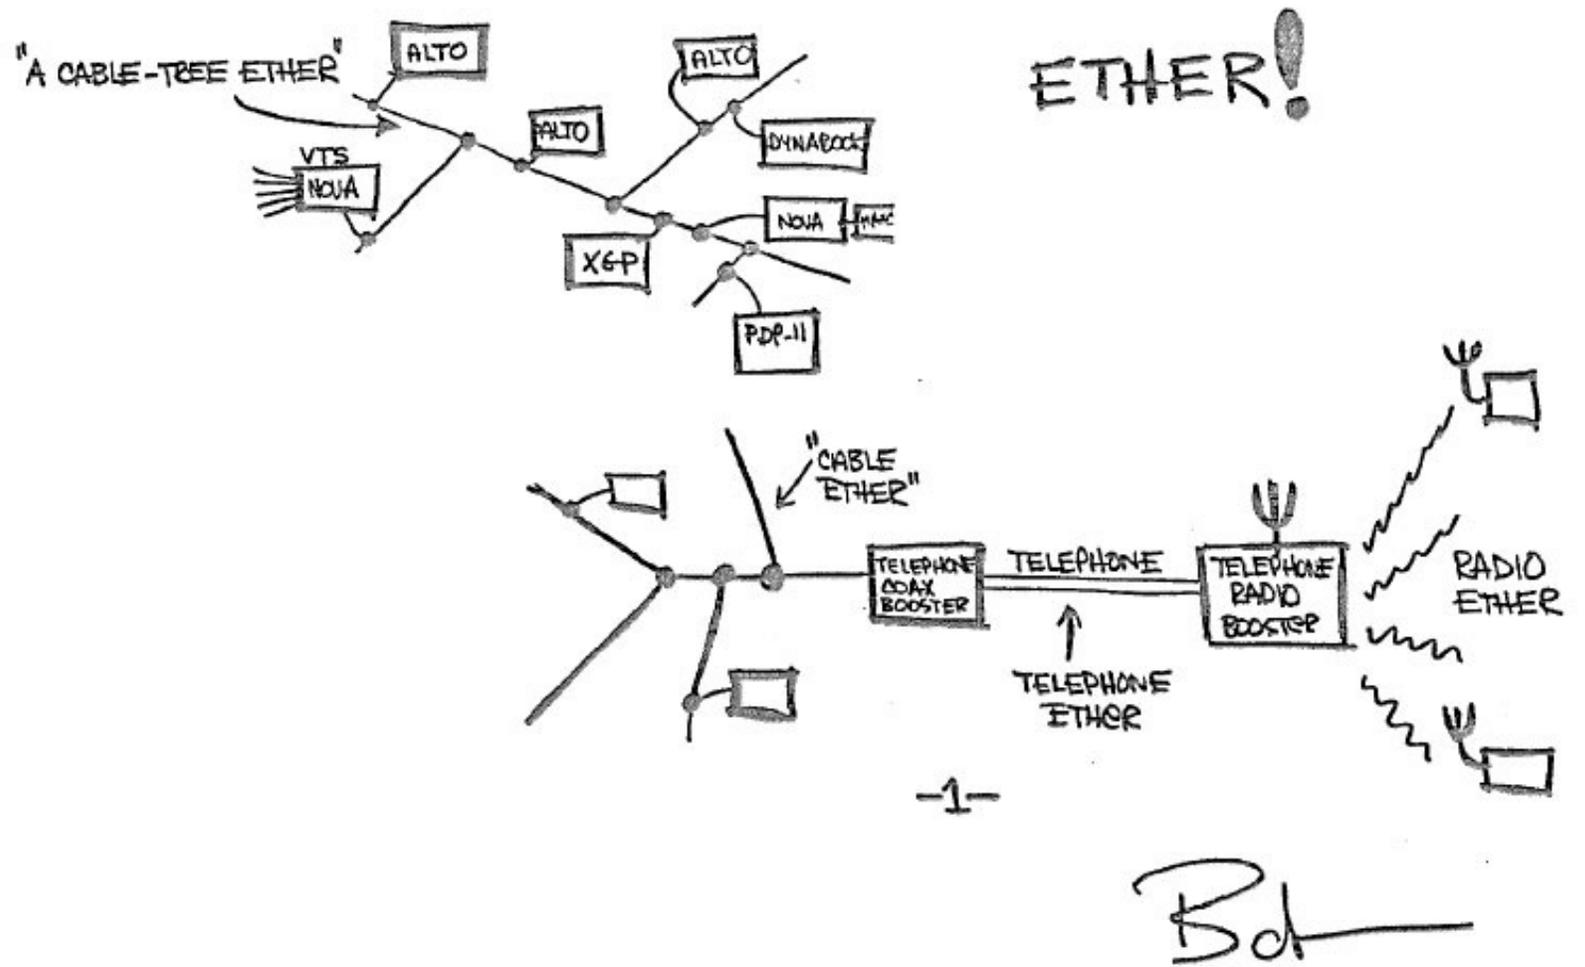
A discussion on how networking
in support of data intensive
research is not at all the same as
networking for general use.



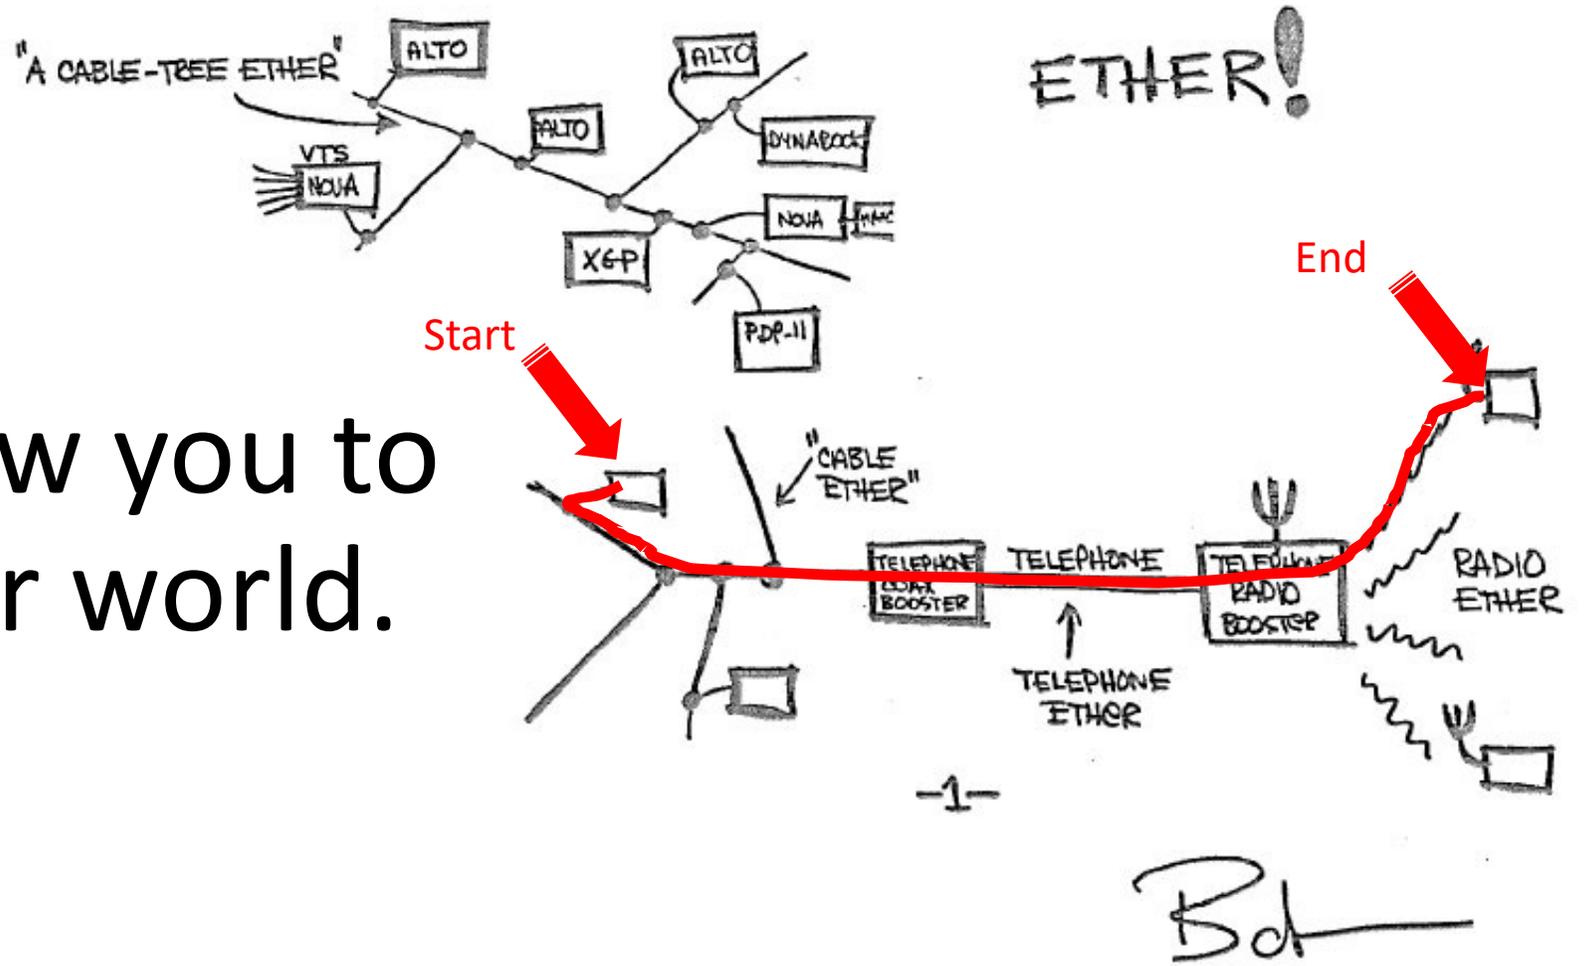
Will this make me A Network Engineer?



No.



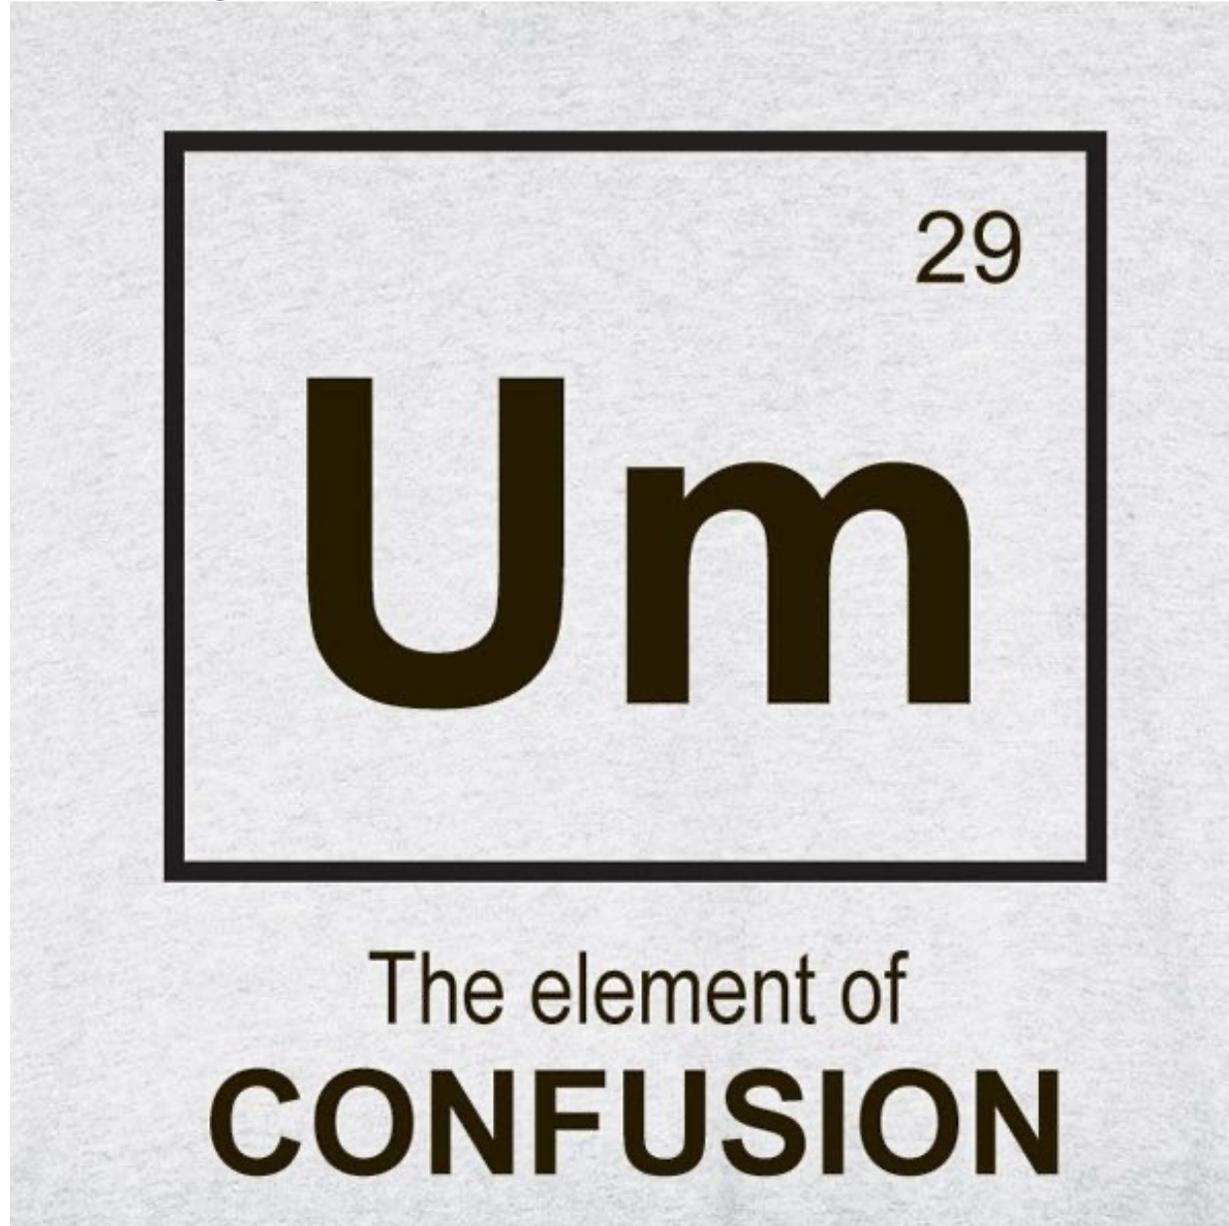
But it will allow you to navigate their world.

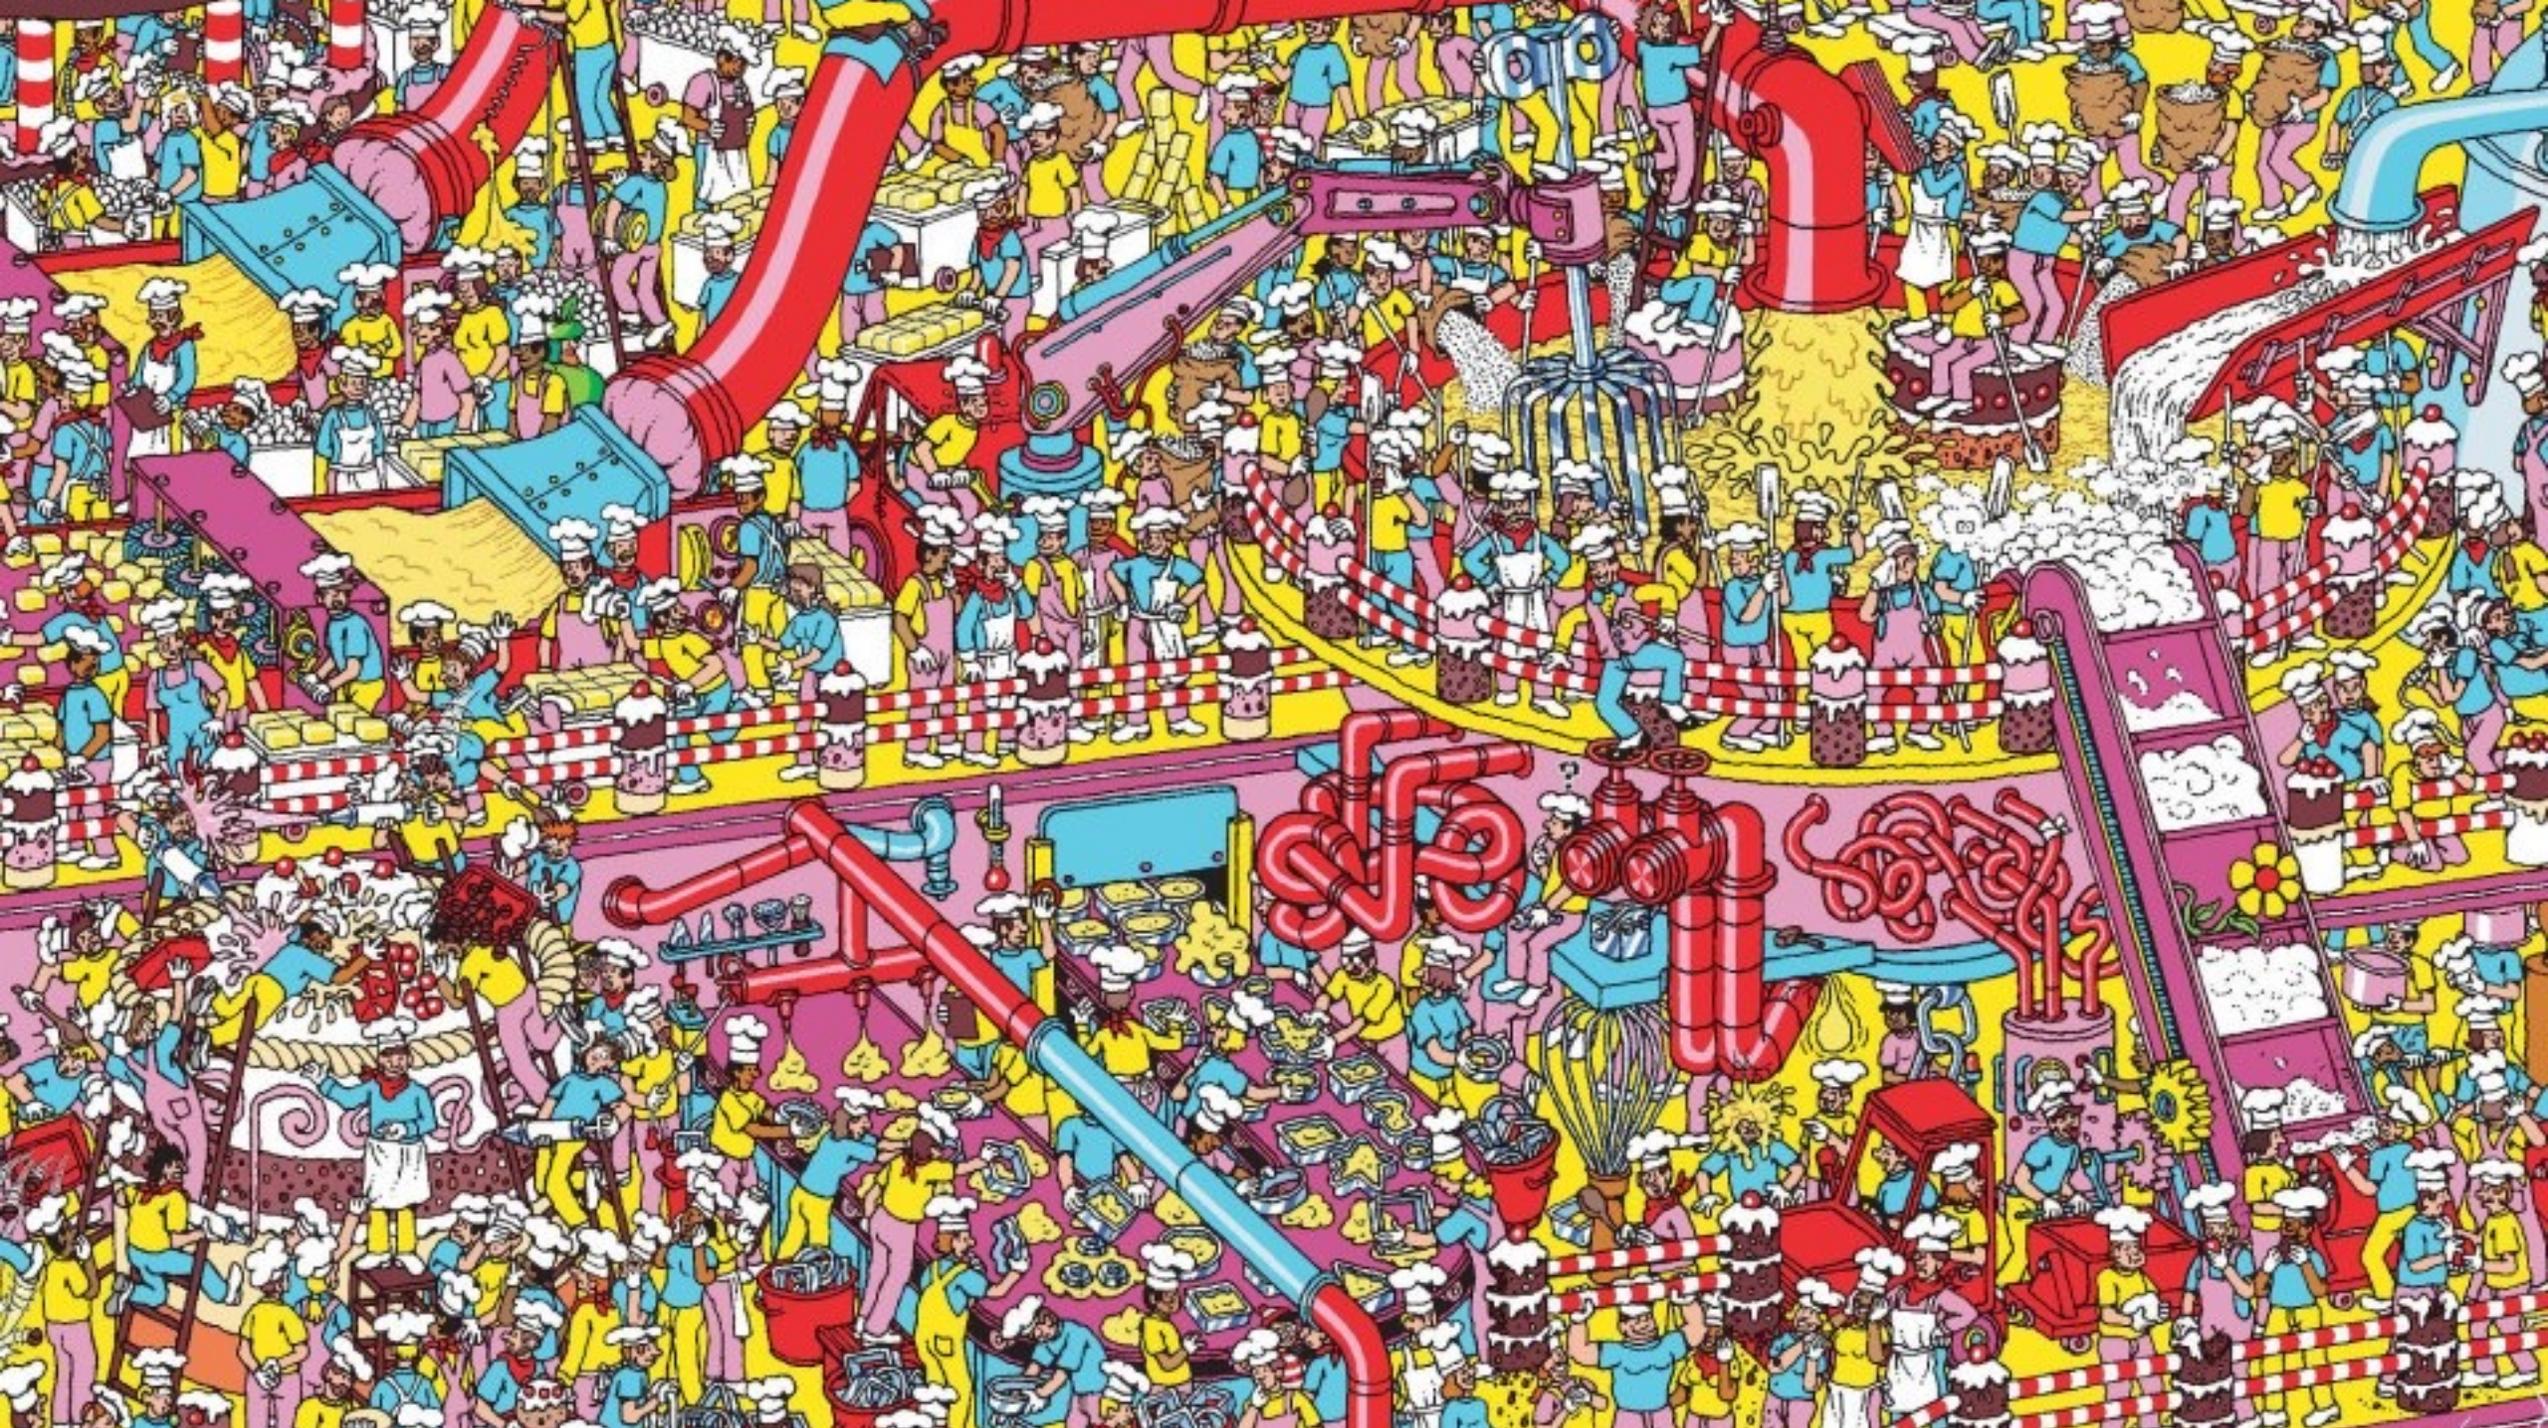


Maybe even understand
them...



But why do I need to know this?



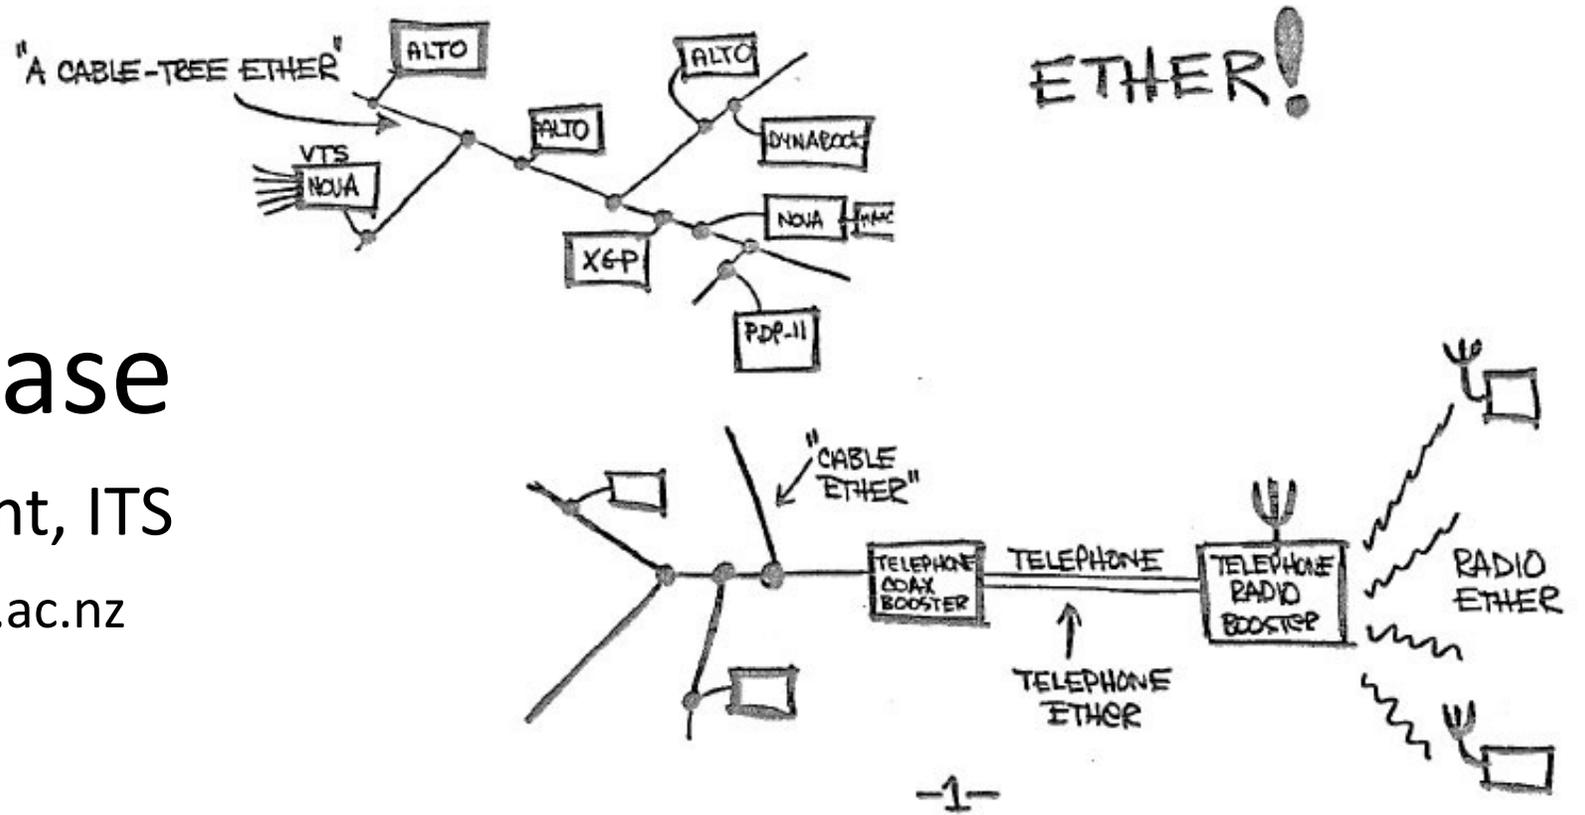


Wallace Chase

Head of Department, ITS

wallace.chase@otago.ac.nz

@bmtfr



UNIVERSITY
of
OTAGO

Te Whare Wānanga o Otāgo
NEW ZEALAND

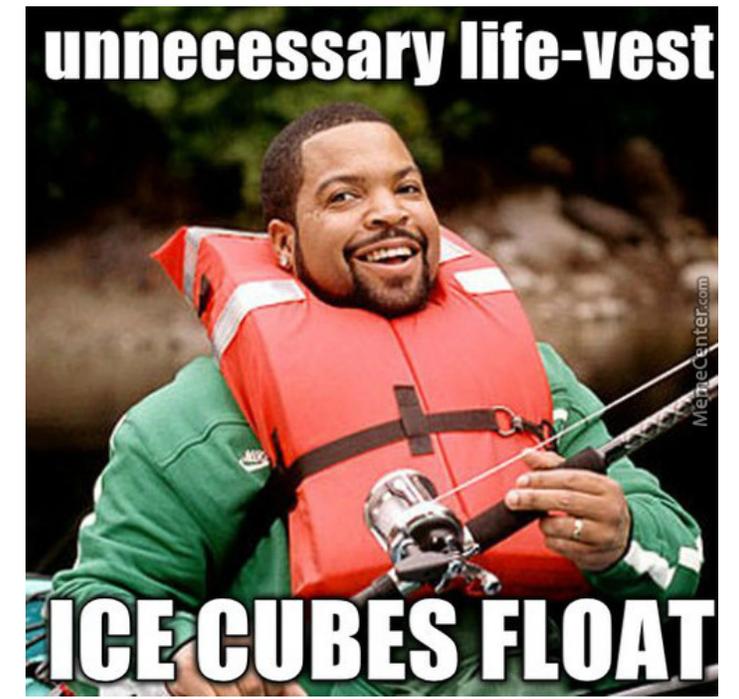


Bch



So, what is
the
difference in
the data?







Research data

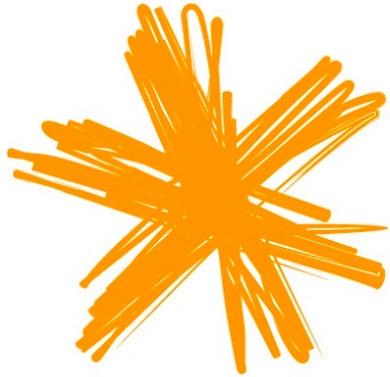


Research data

Who provides the transport?



Commodity/commercial Networks



Spark^{nz}



orcon

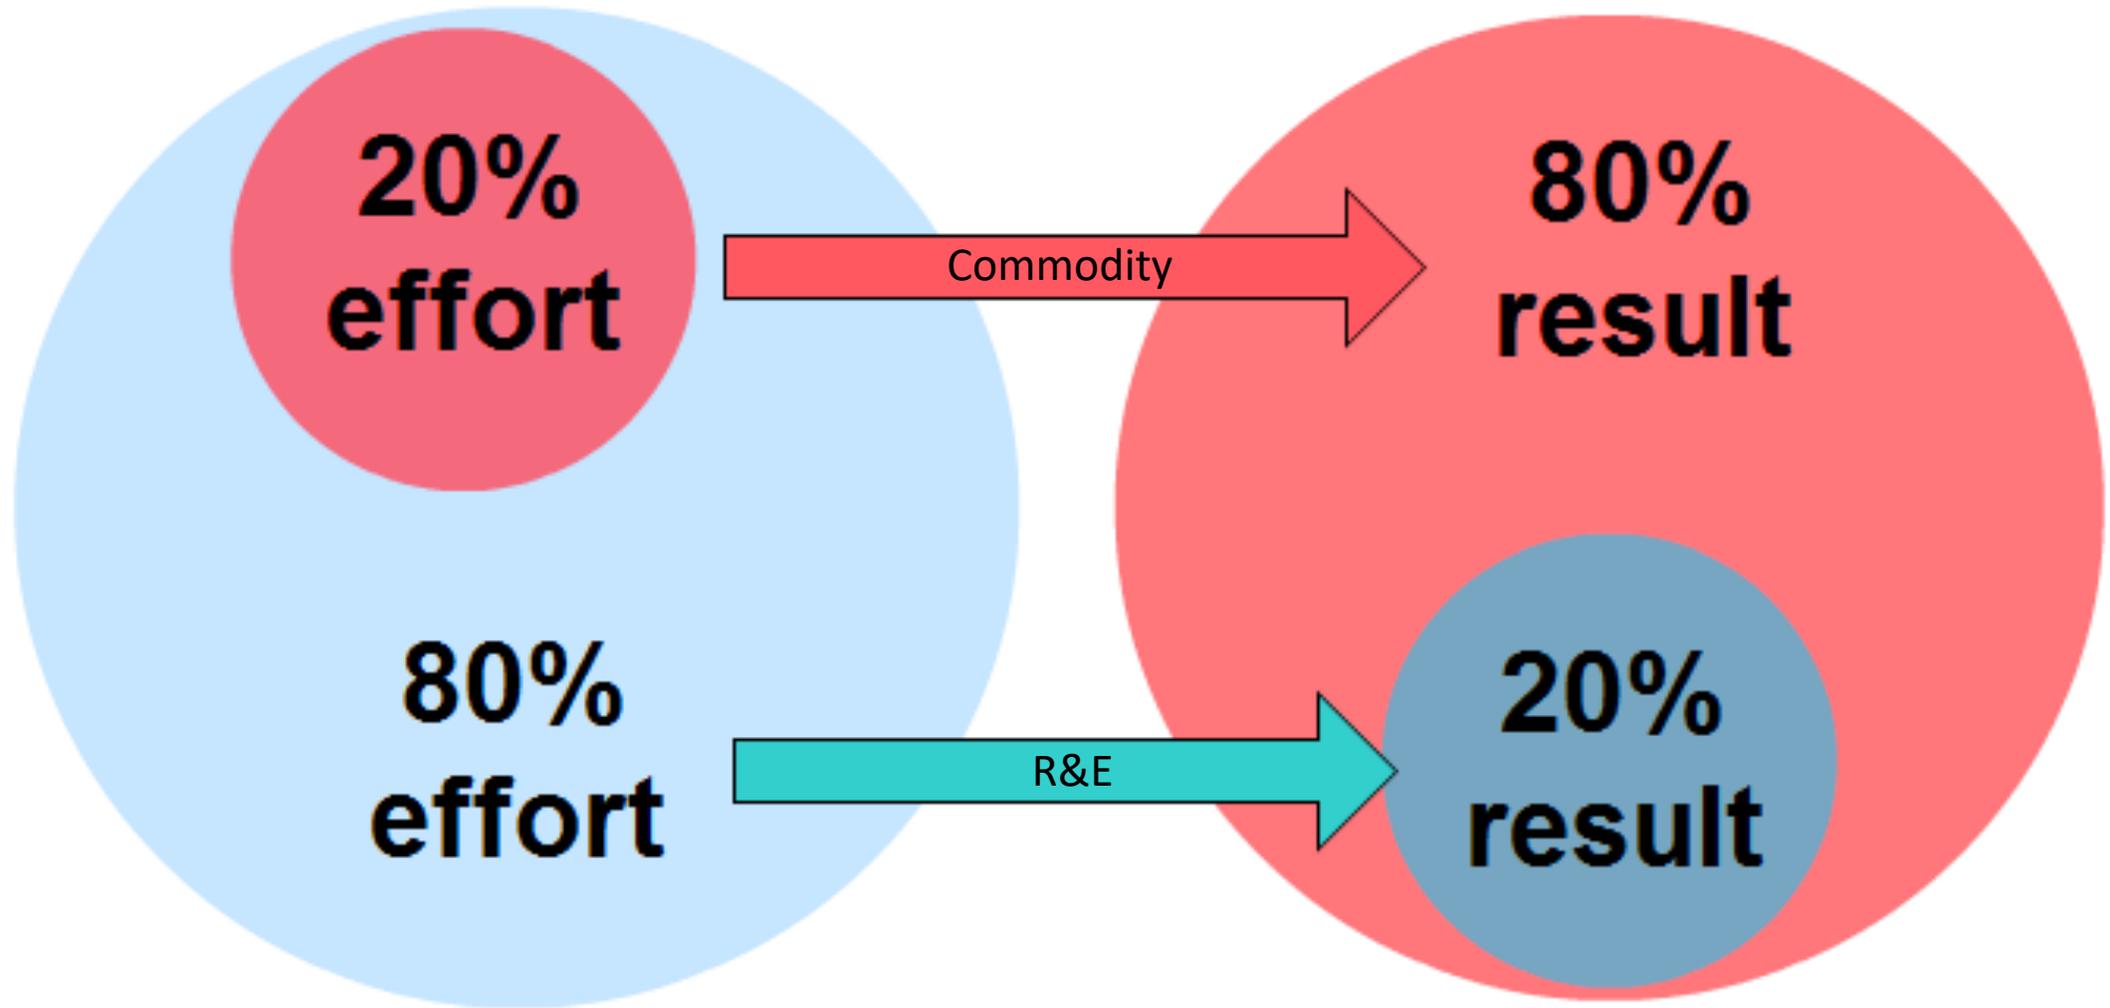


TeliaSonera

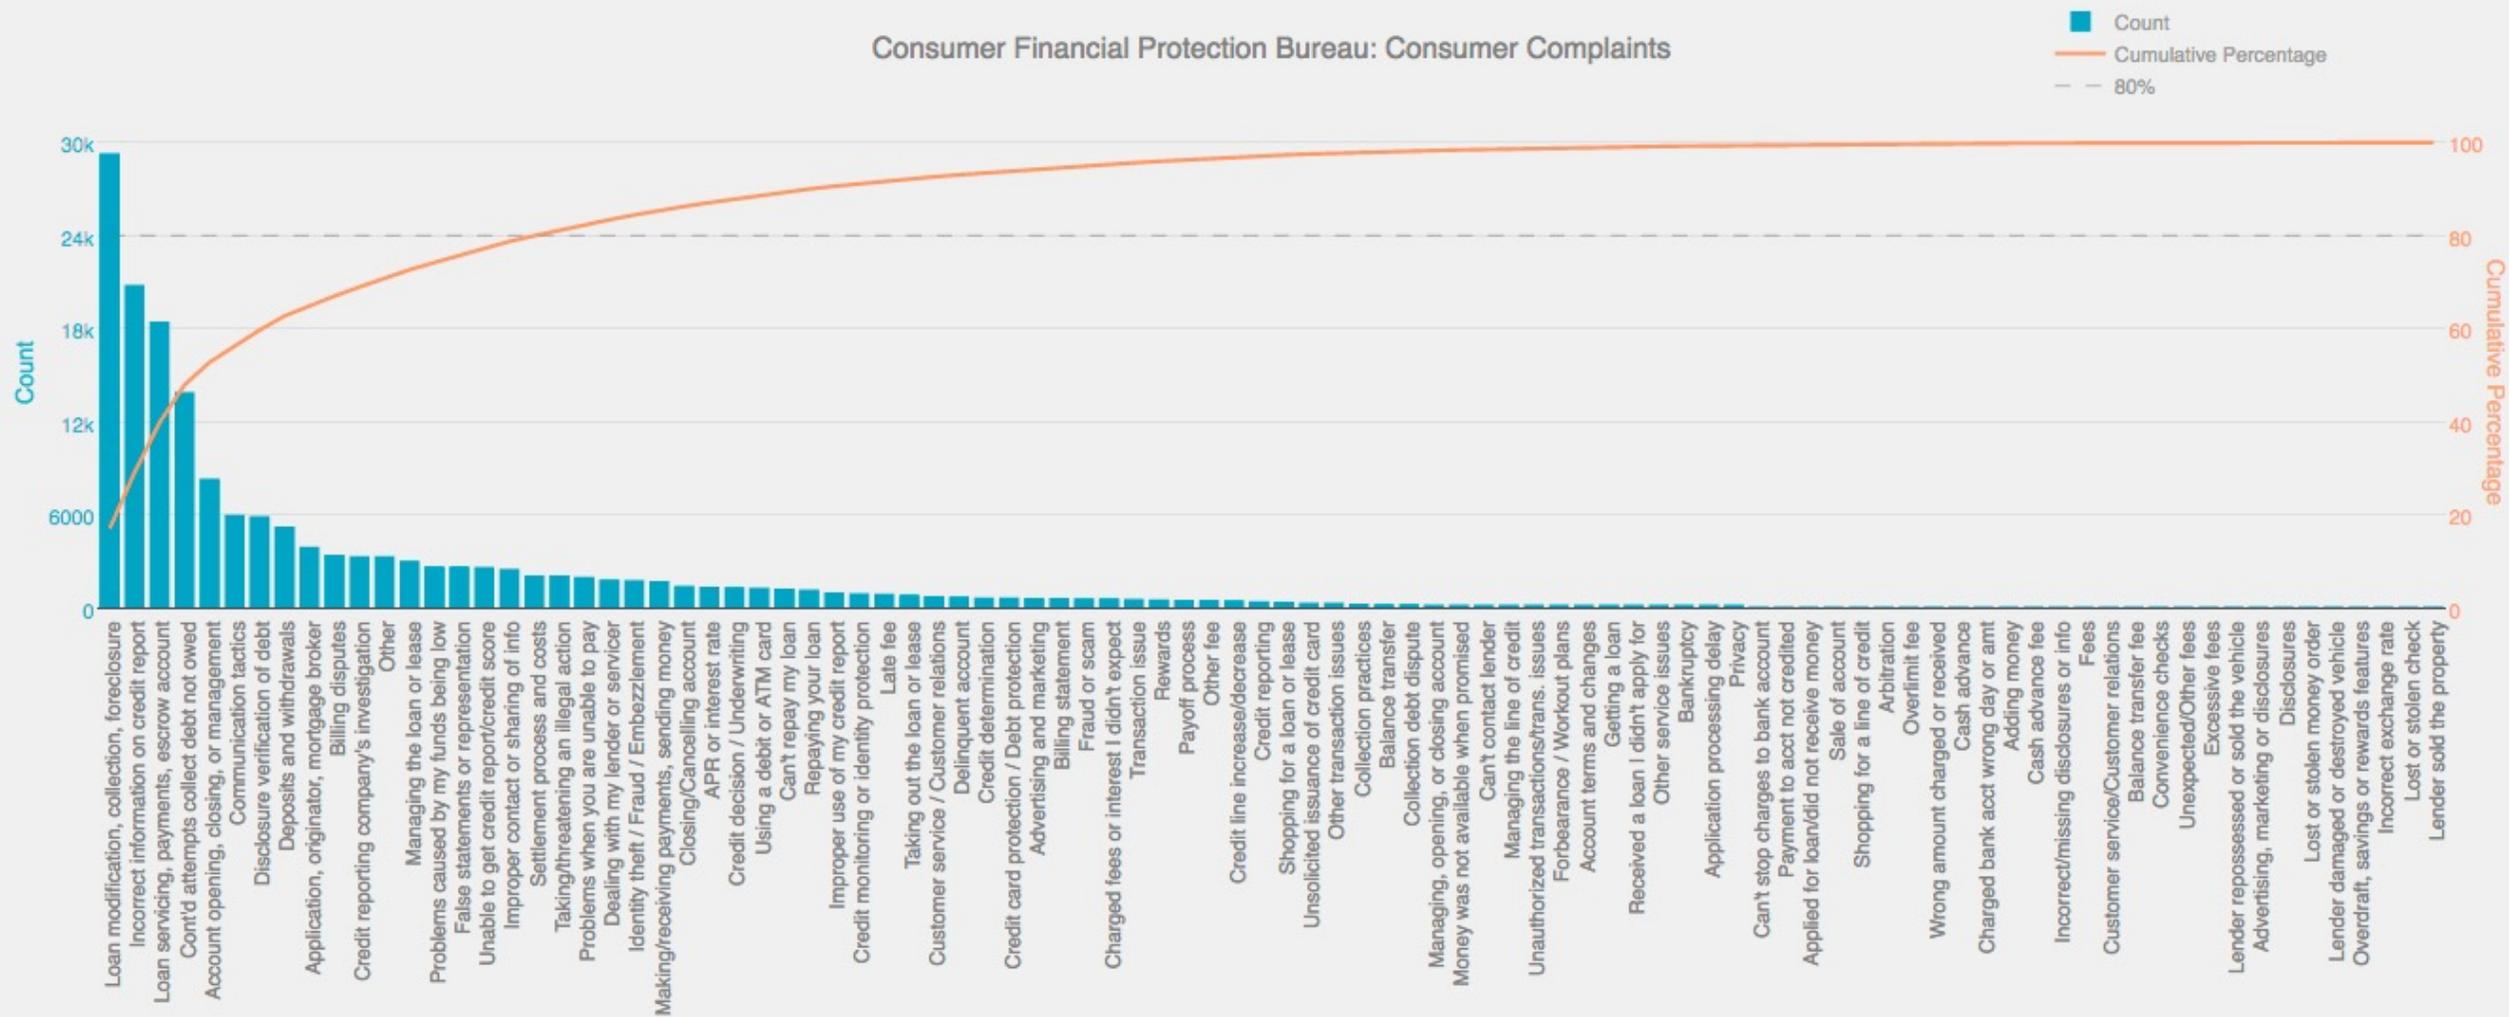


at&t

Effort

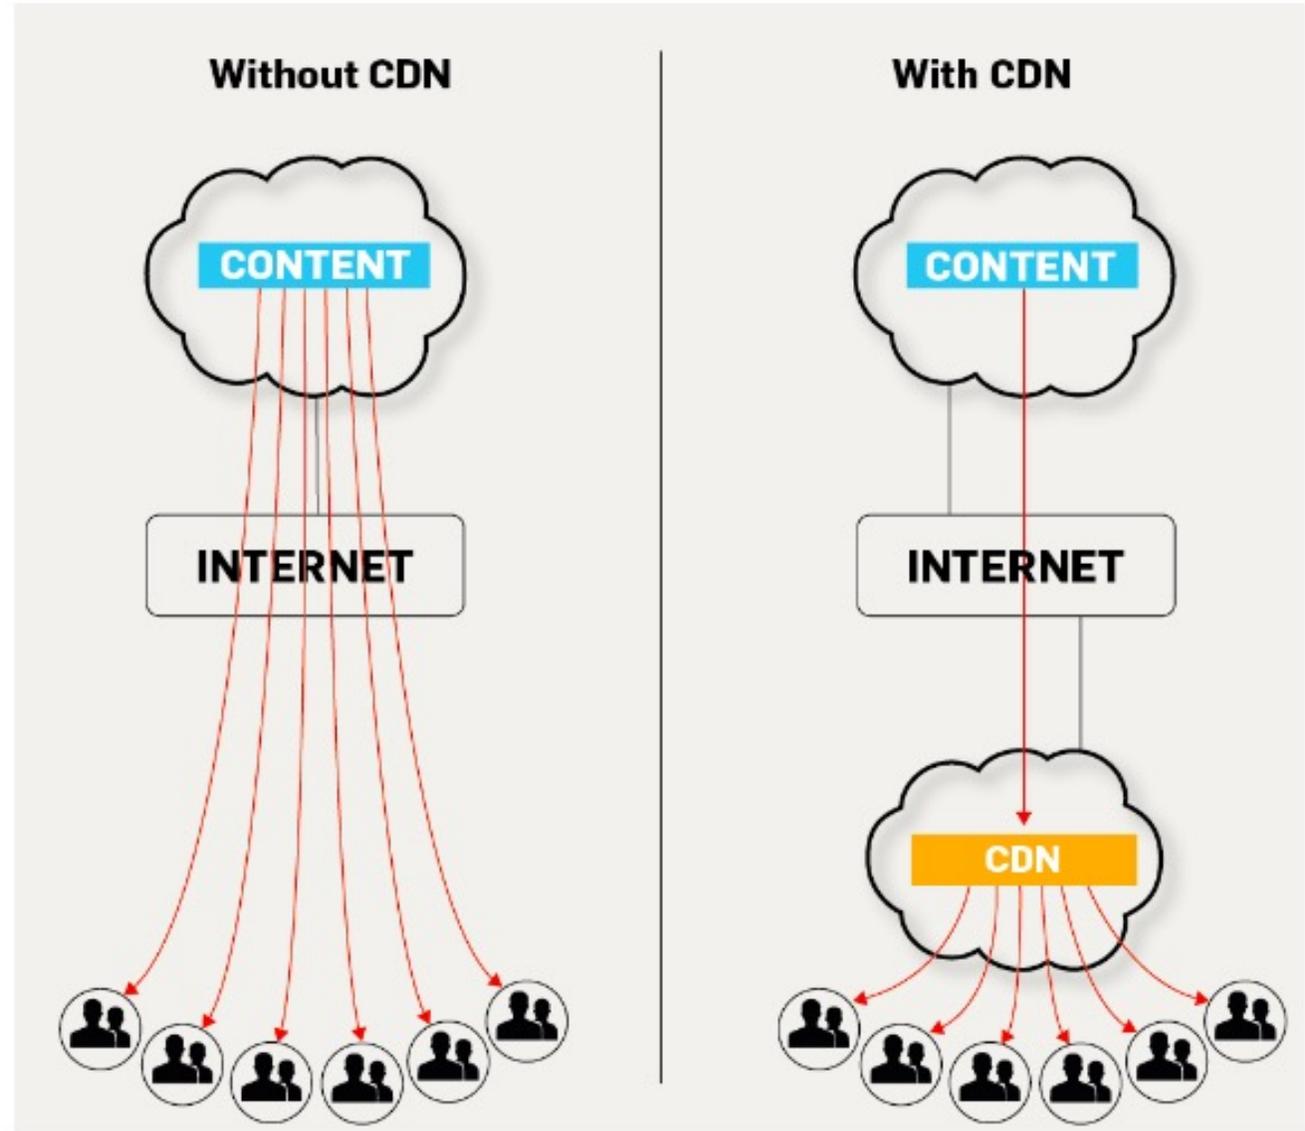


Consumer Financial Protection Bureau: Consumer Complaints



Content Delivery Network

- Globally distributed network of web servers or Points of Presence (PoP) whose purpose is to provide faster content delivery.
- Content is replicated and stored throughout the CDN so the user can access the data that is stored at a location that is geographically closest to the user.



An example of traffic flows...

Volumes Transferred



2.52PB

Total Traffic



462TB

International Commodity



734TB

Domestic Commodity



30.6TB

International NREN



1.06PB

Domestic NREN

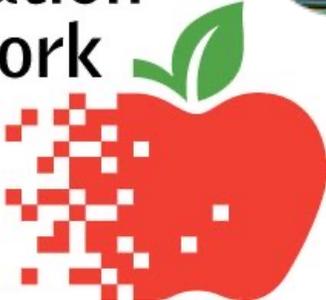


DEPTH
GAUGE

State/regional networks aka "your ISP"



K-20
Education
Network





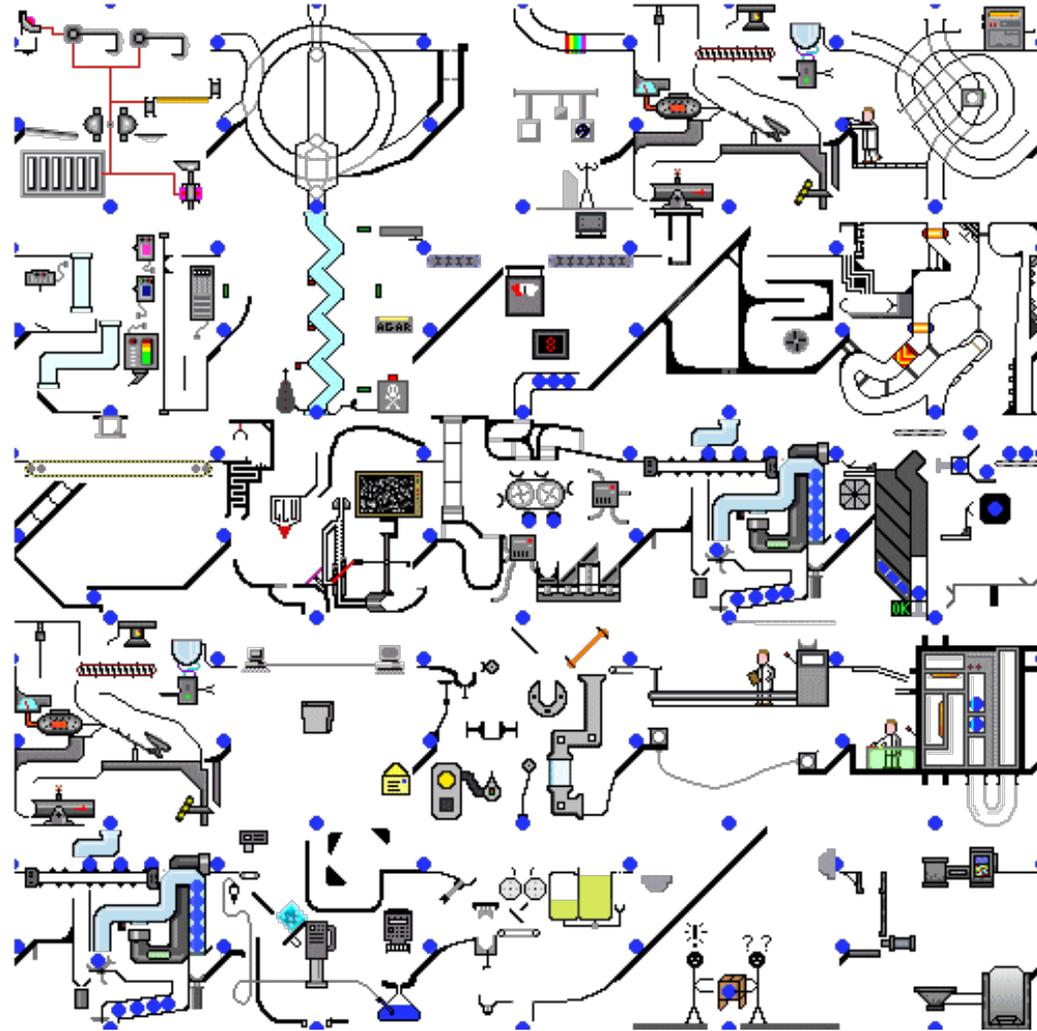
But they go
about it in
different ways...

They both eat
the leaves...

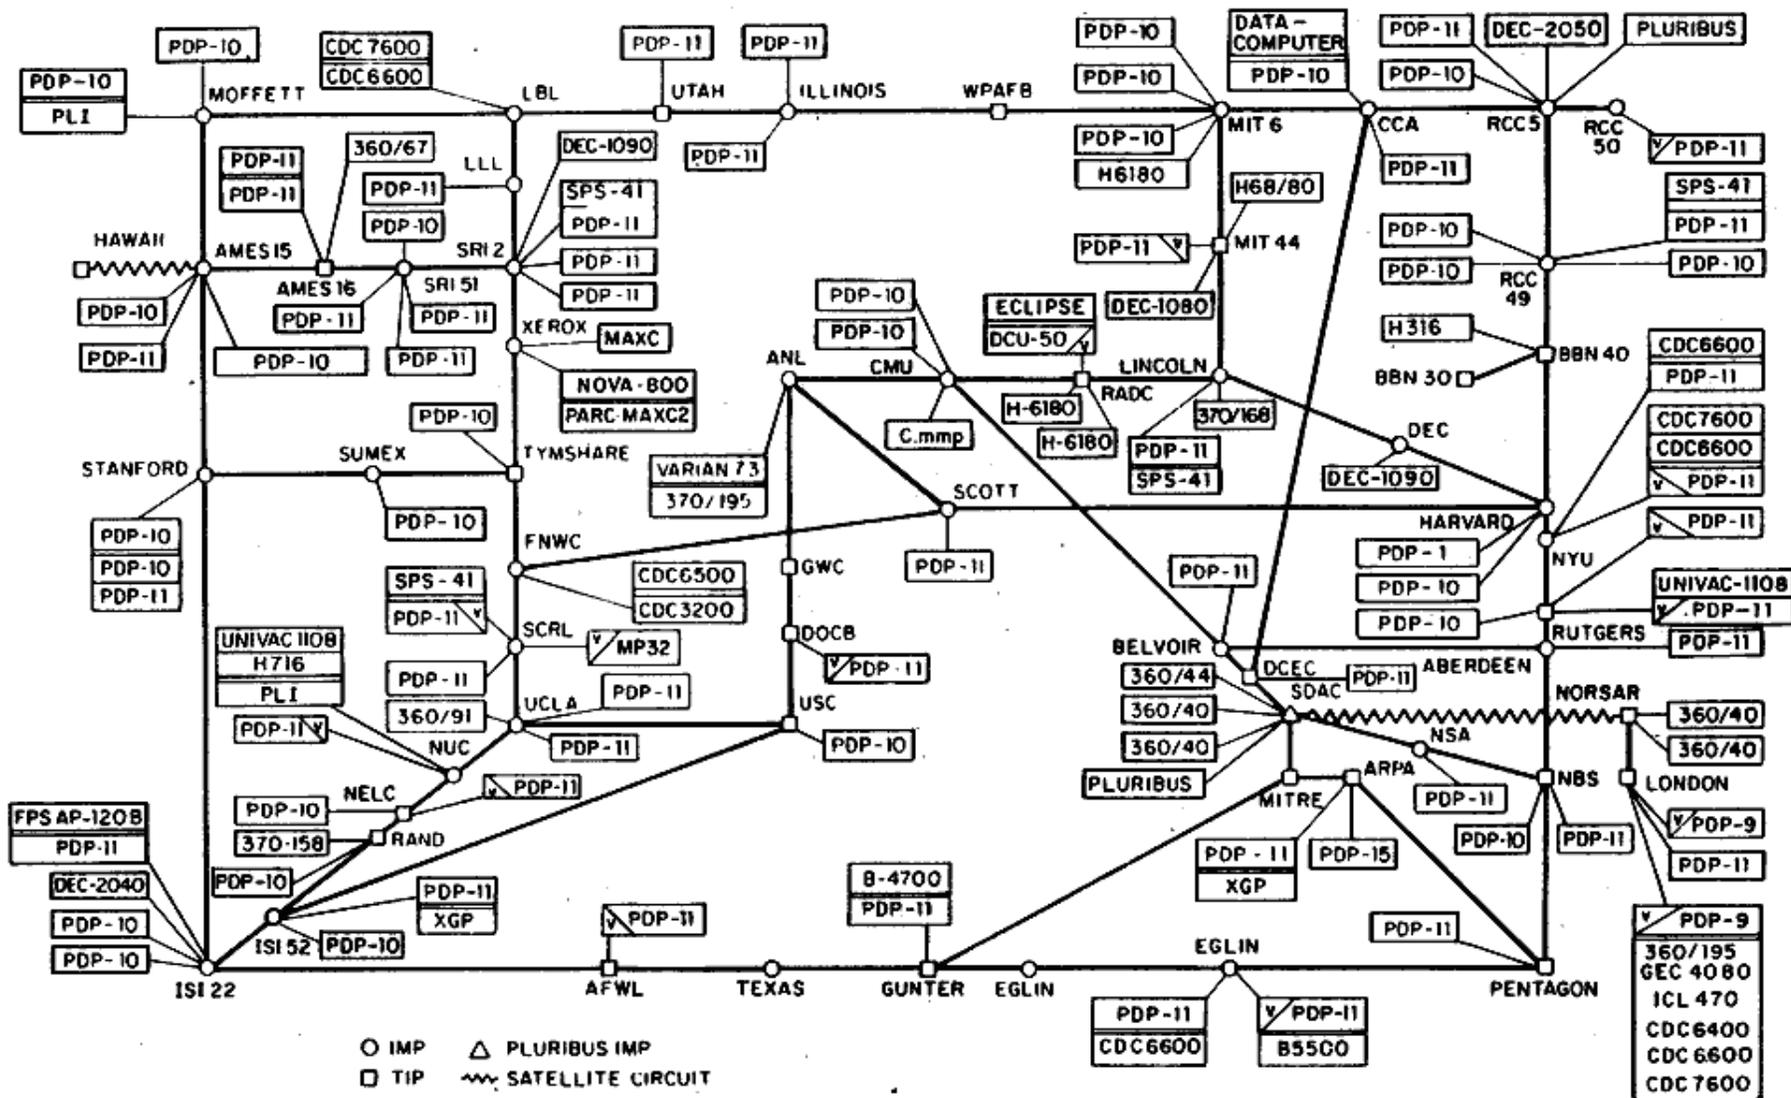




Now to the technical reasons...



ARPANET LOGICAL MAP, MARCH 1977



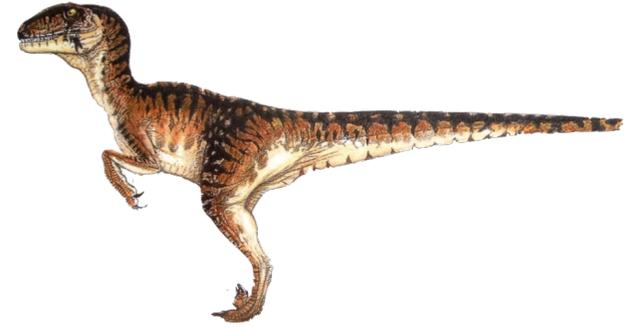
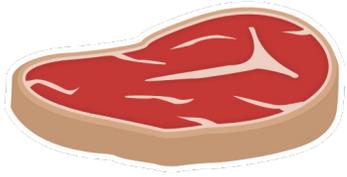
(PLEASE NOTE THAT WHILE THIS MAP SHOWS THE HOST POPULATION OF THE NETWORK ACCORDING TO THE BEST INFORMATION OBTAINABLE, NO CLAIM CAN BE MADE FOR ITS ACCURACY)

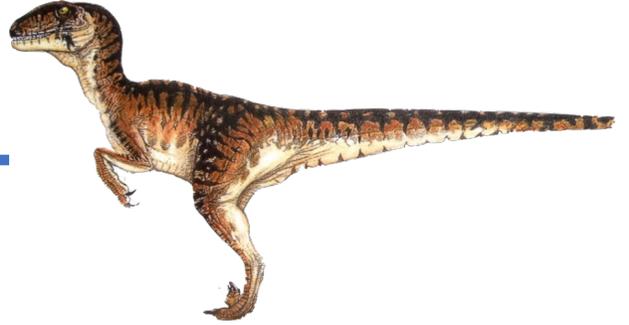
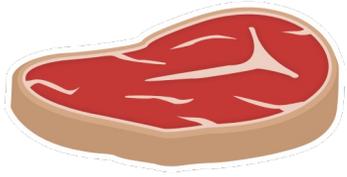
NAMES SHOWN ARE IMP NAMES, NOT (NECESSARILY) HOST NAMES

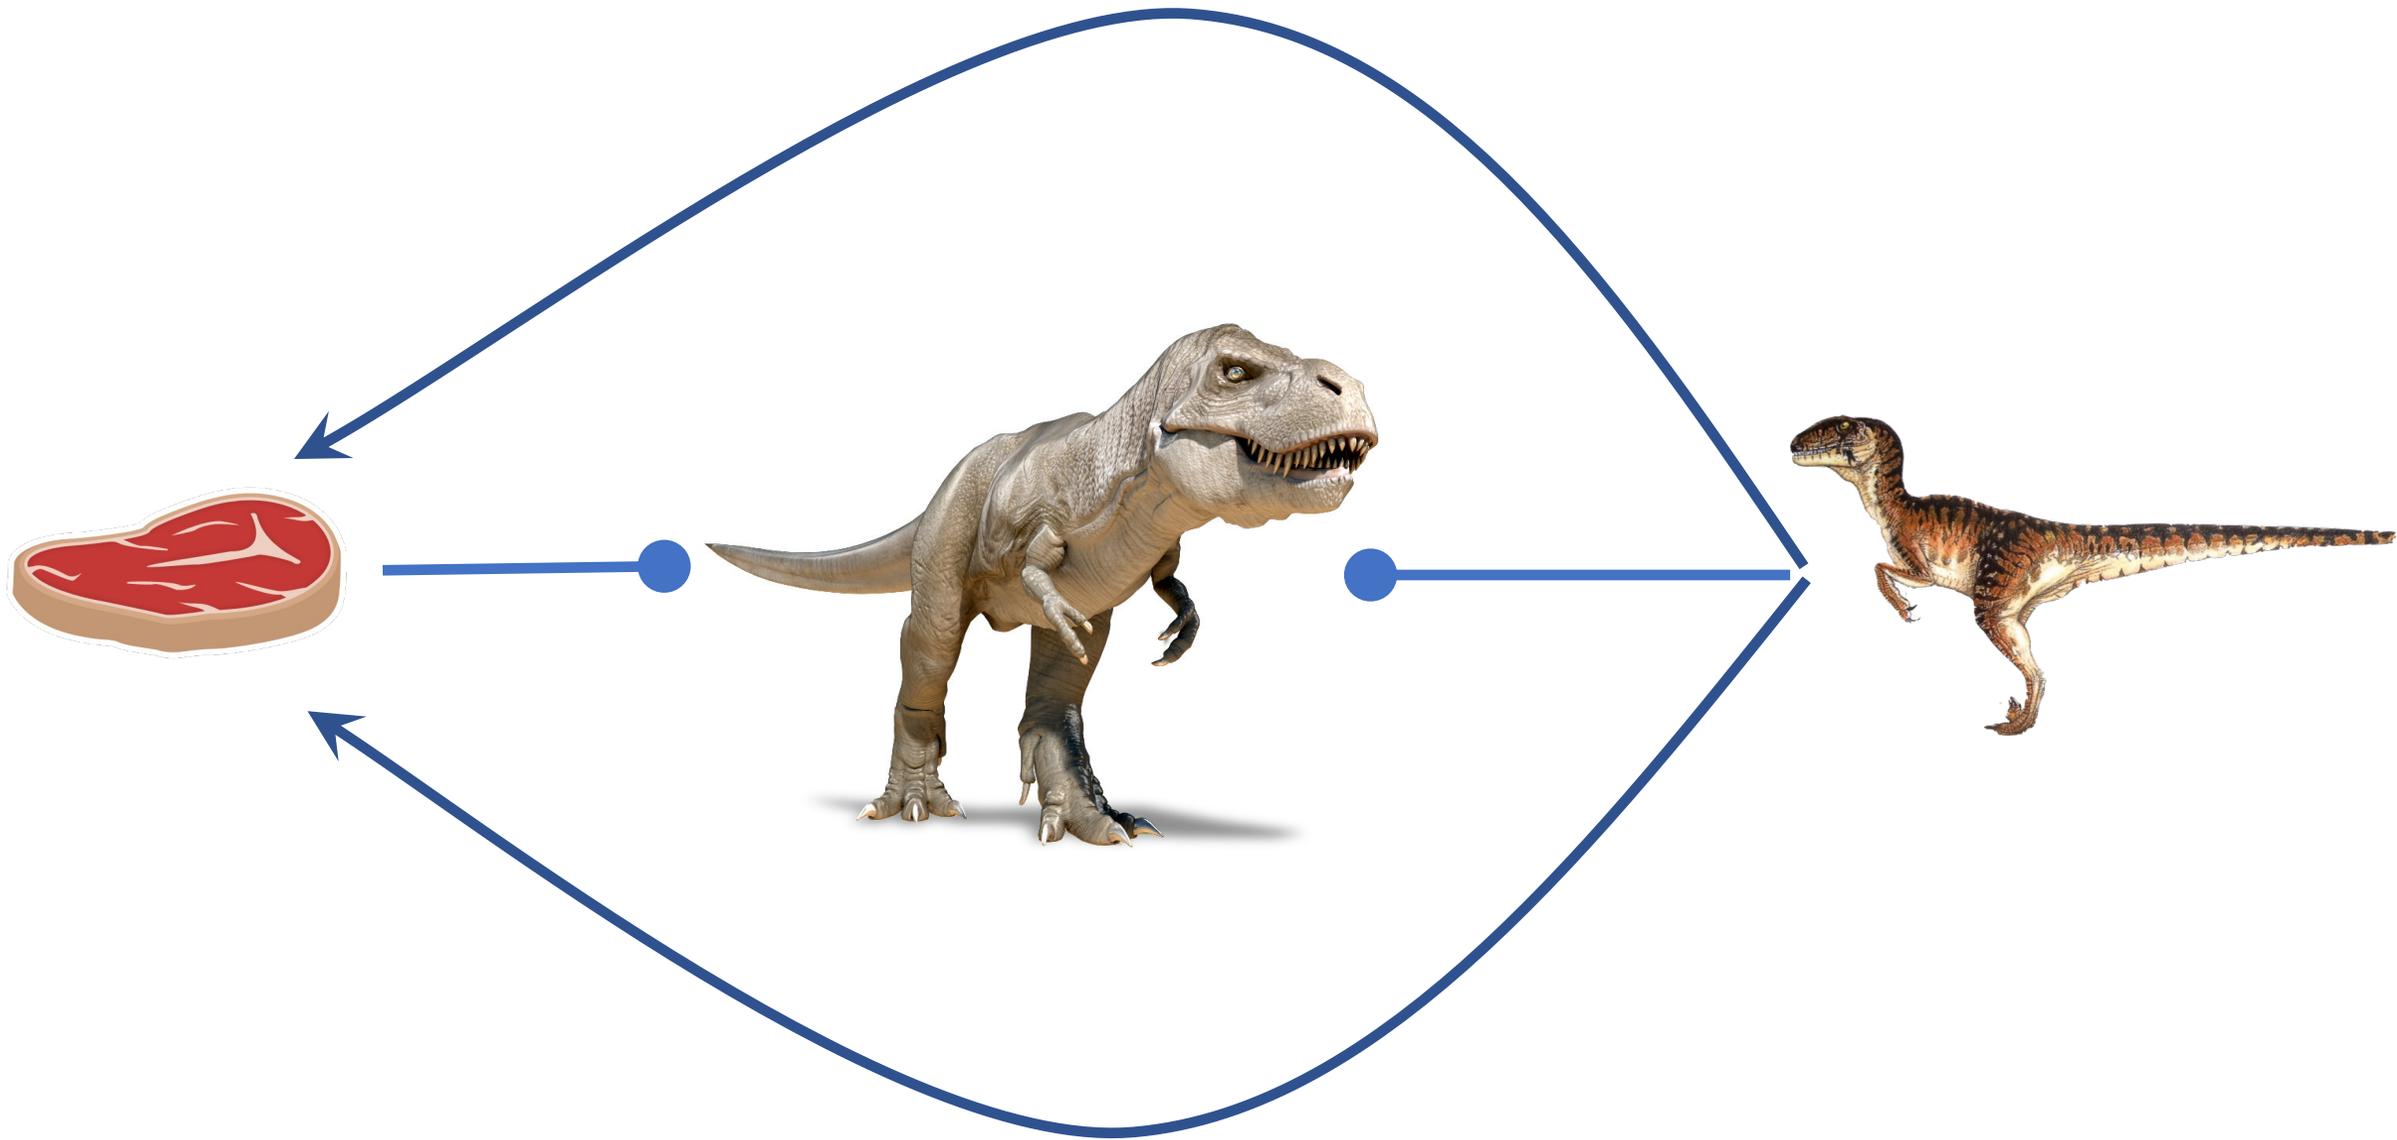
The goal was to exploit new computer technologies to meet the needs of military command and control against nuclear threats, achieve survivable control of US nuclear forces, and improve military tactical and management decision making.

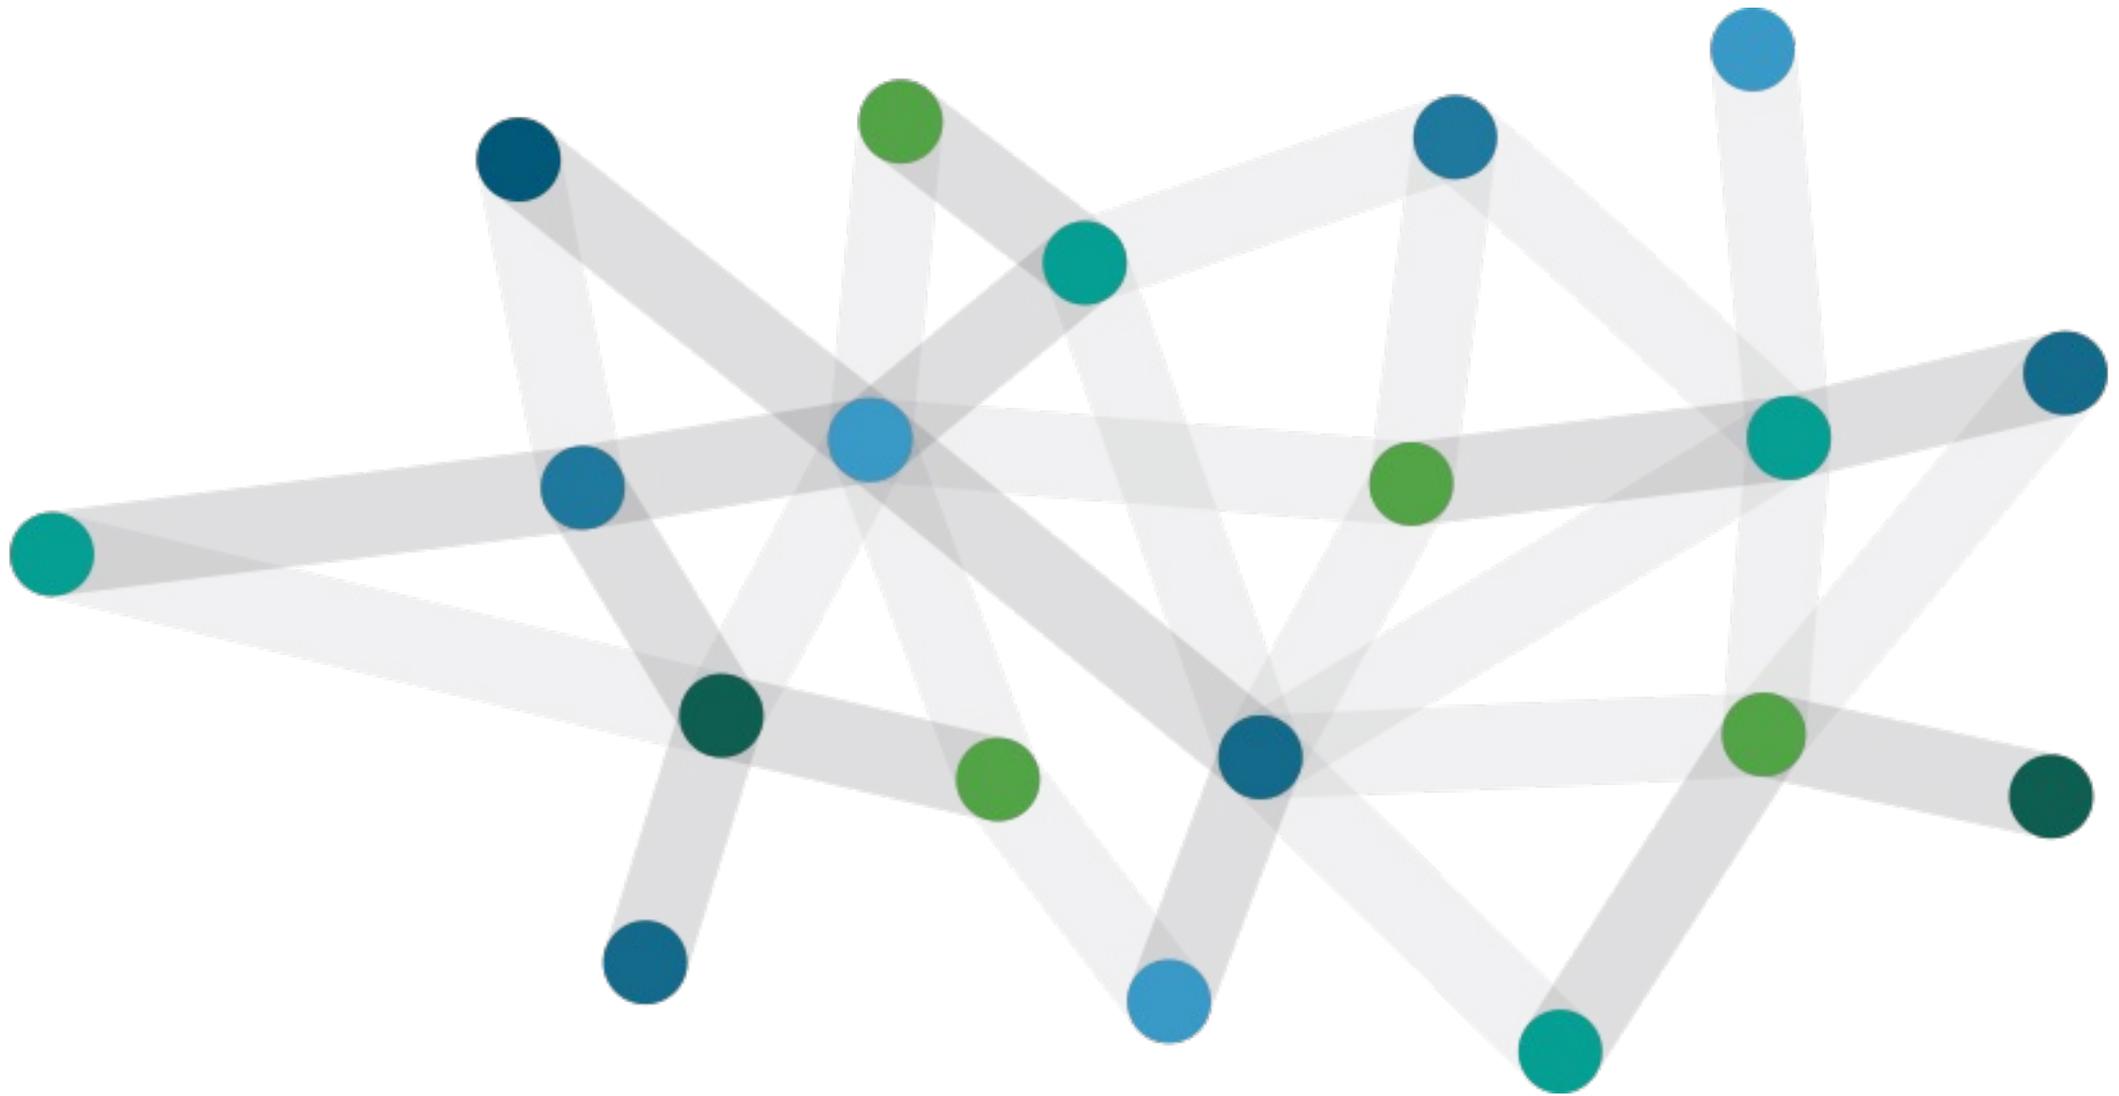
Stephen J. Lukasik



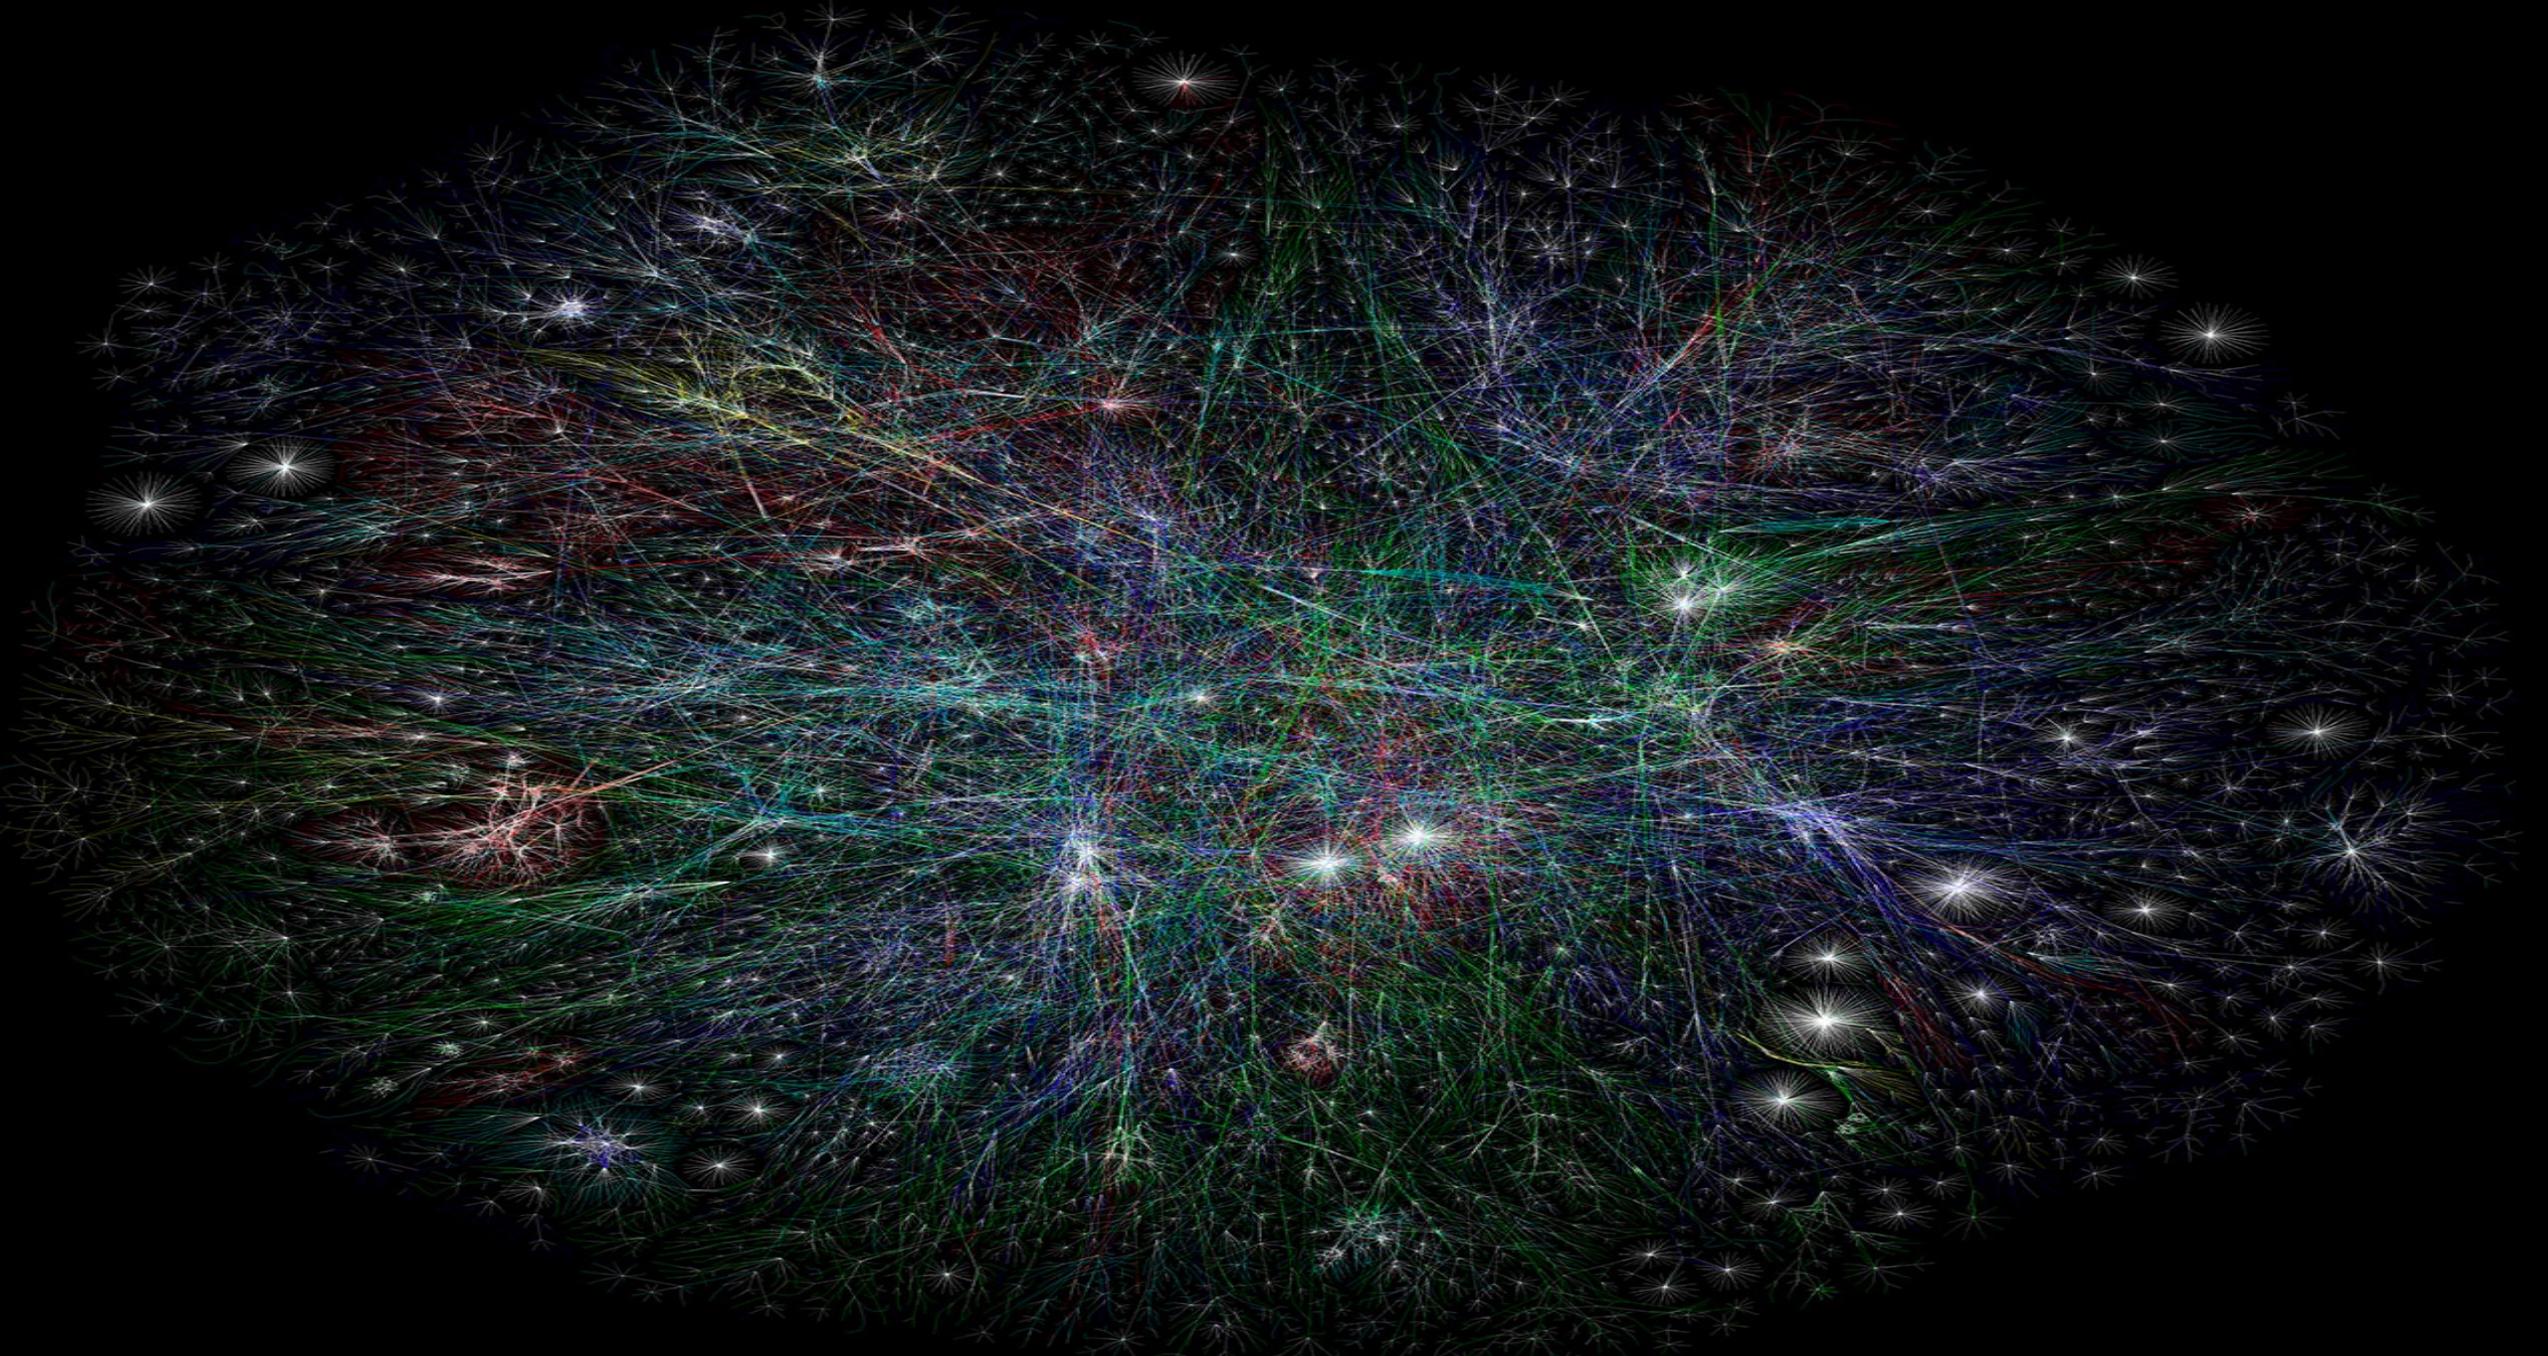








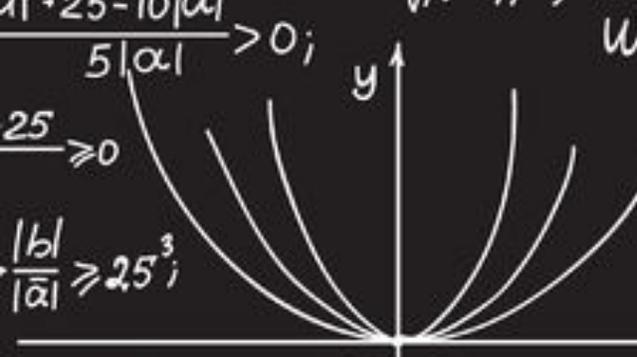
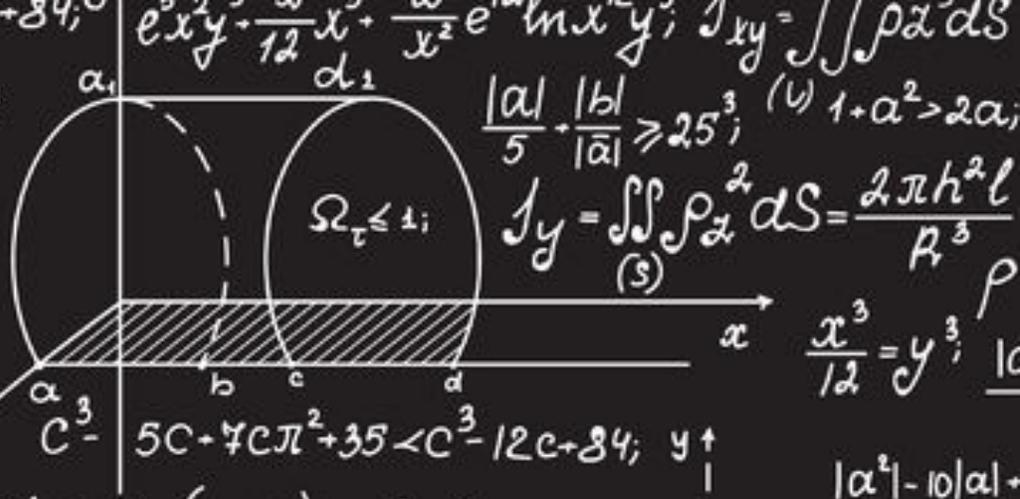
Autonomous systems, interconnecting using standardized protocols...



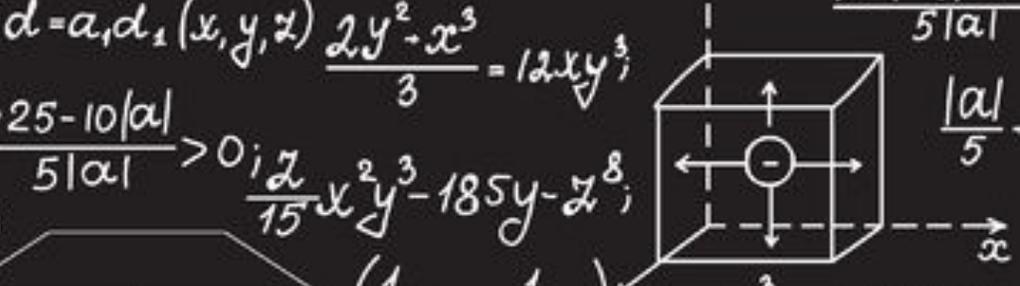
$$\begin{cases} x \cdot y = \frac{x}{12} \Rightarrow \frac{1}{12} = y; \frac{|a| - 10|a| + 25}{5|a|} \geq 0 \Rightarrow \frac{(|a| - 5)^2}{5|a|} \geq 0; \left(\frac{1}{R} - \frac{1}{\sqrt{R^2 - h^2}}\right); 1 + a^2 > 2a; \\ \frac{x^3 - y^{12}}{x} = 0; \frac{(c^2 - 1)^2}{c} \geq 0 \Rightarrow \frac{(c^2 - 1)^2}{c} \geq 0 \end{cases}$$

$$\frac{2y^2 - x^3}{3} = 12xy^3; \frac{c^2 + \frac{1}{2}\pi^3}{2} \geq 0; \left(\frac{1}{R} - \frac{1}{\sqrt{R^2 - h^2}}\right);$$

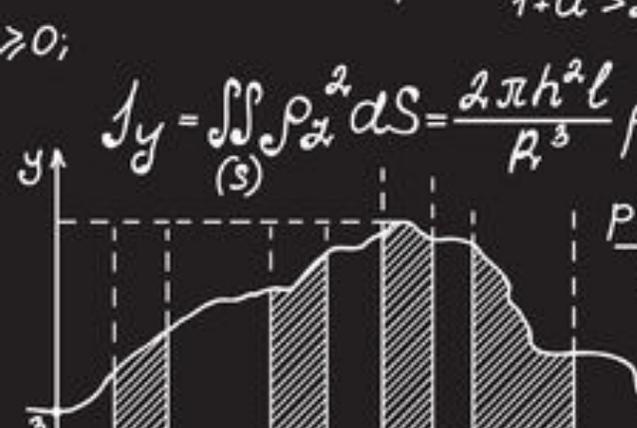
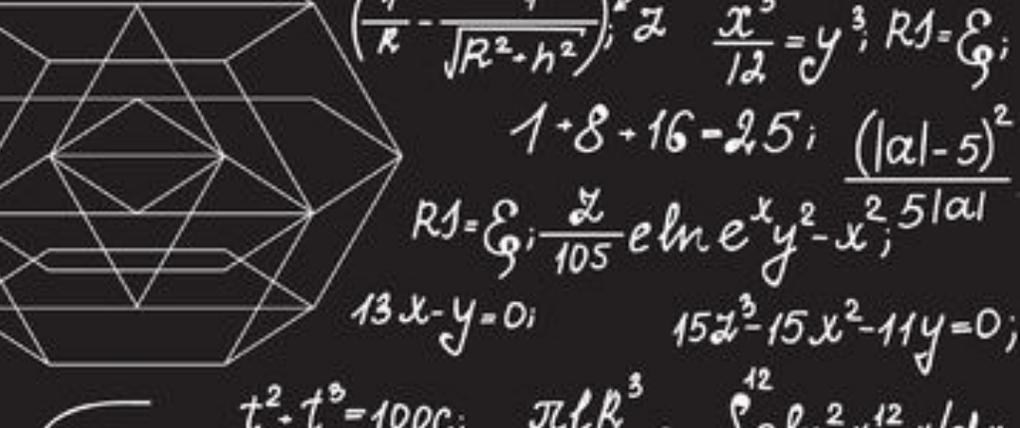
$$F_2 = \int_0^{25} \int_0^h \frac{2R dz d\theta}{(R^2 - z^2)^{\frac{3}{2}}} = 2\pi R \int_0^h \frac{dz}{\sqrt{R^2 - z^2}}; \frac{m^2 + s^3}{\pi}$$



$$W = \int_0^{2\pi} \int_0^1 \frac{R dz d\theta}{\sqrt{R^2 - z^2}}; \int_{yz} = \int_{2x} = \frac{\pi l R^3}{4} \rho$$



$$\frac{3y^2 - 105z^3 - x^3}{15} = 0 = z; \int_0^R r^3 dr = \frac{\pi h^2 l R}{2} \rho;$$



$$\begin{cases} x + y = \frac{x^2}{12} \\ \frac{x^3 - y^{12}}{x} = 0; \end{cases} \frac{p^2 - m^3}{z} - 100y = x;$$

$$\int y dy \quad \int x dx \quad \int x dy \quad \int y dx$$

$$I = \frac{V_1 - V_0}{\rho}; \quad R(I, \frac{dI}{dt}) = \oint (F_1 + F_2) dl; \quad S^{15} \quad 3y^2 - 105z^3 - x^3$$



Clarke's Third Law: Any sufficiently advanced technology is indistinguishable from magic.

(Arthur C. Clarke)

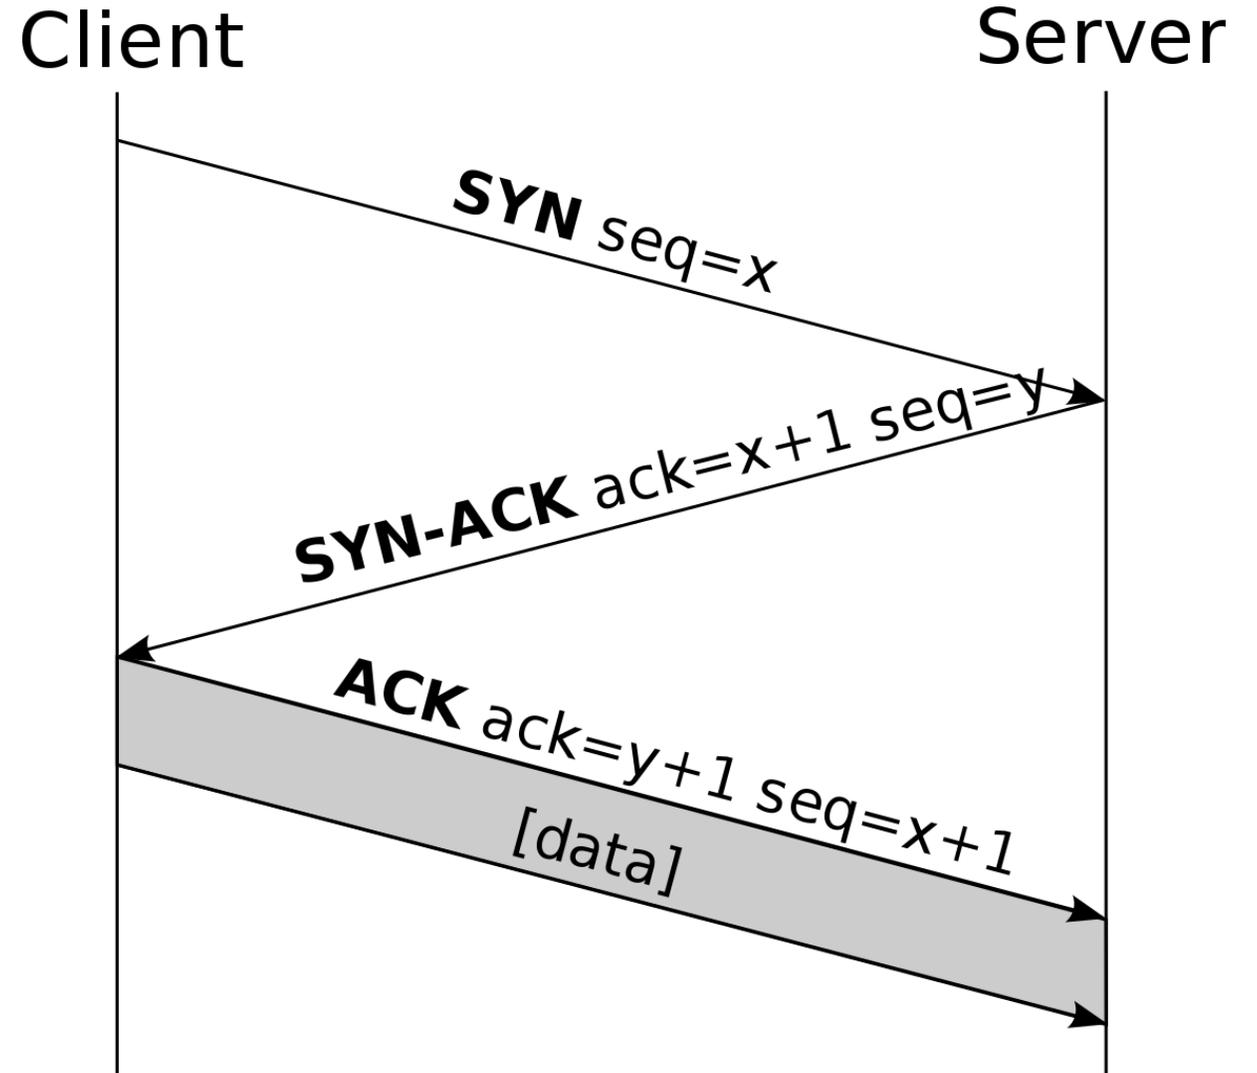
izquotes.com

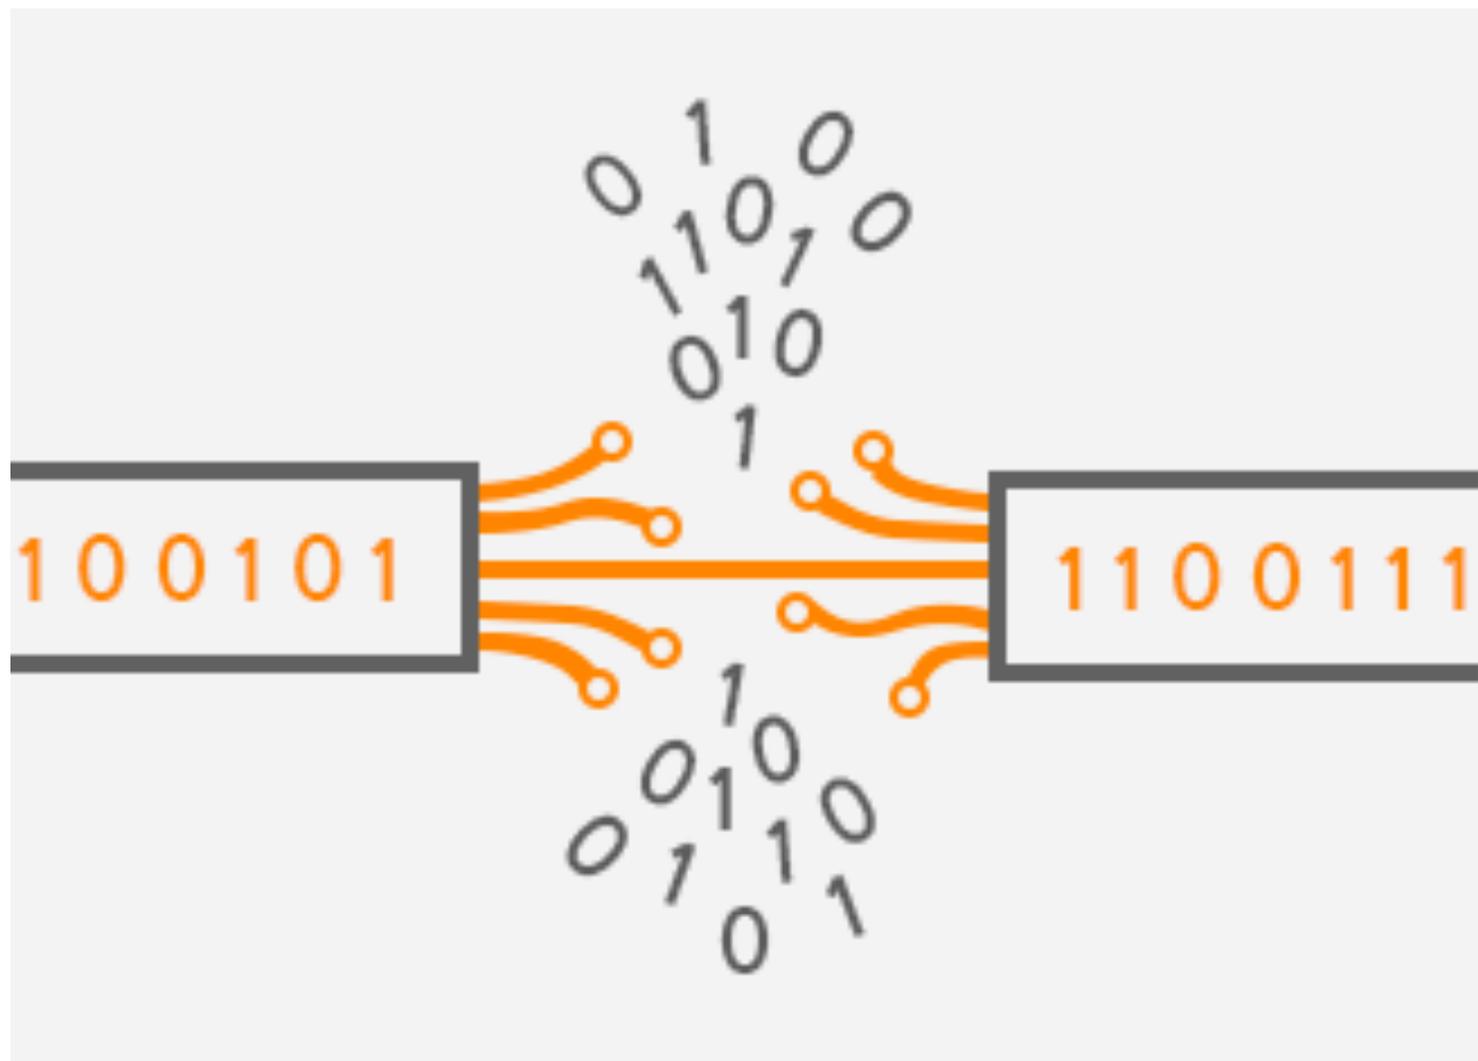
Transmission Control Protocol (TCP)

There are three major factors that affect TCP performance (there are others, but these are the Big Three):

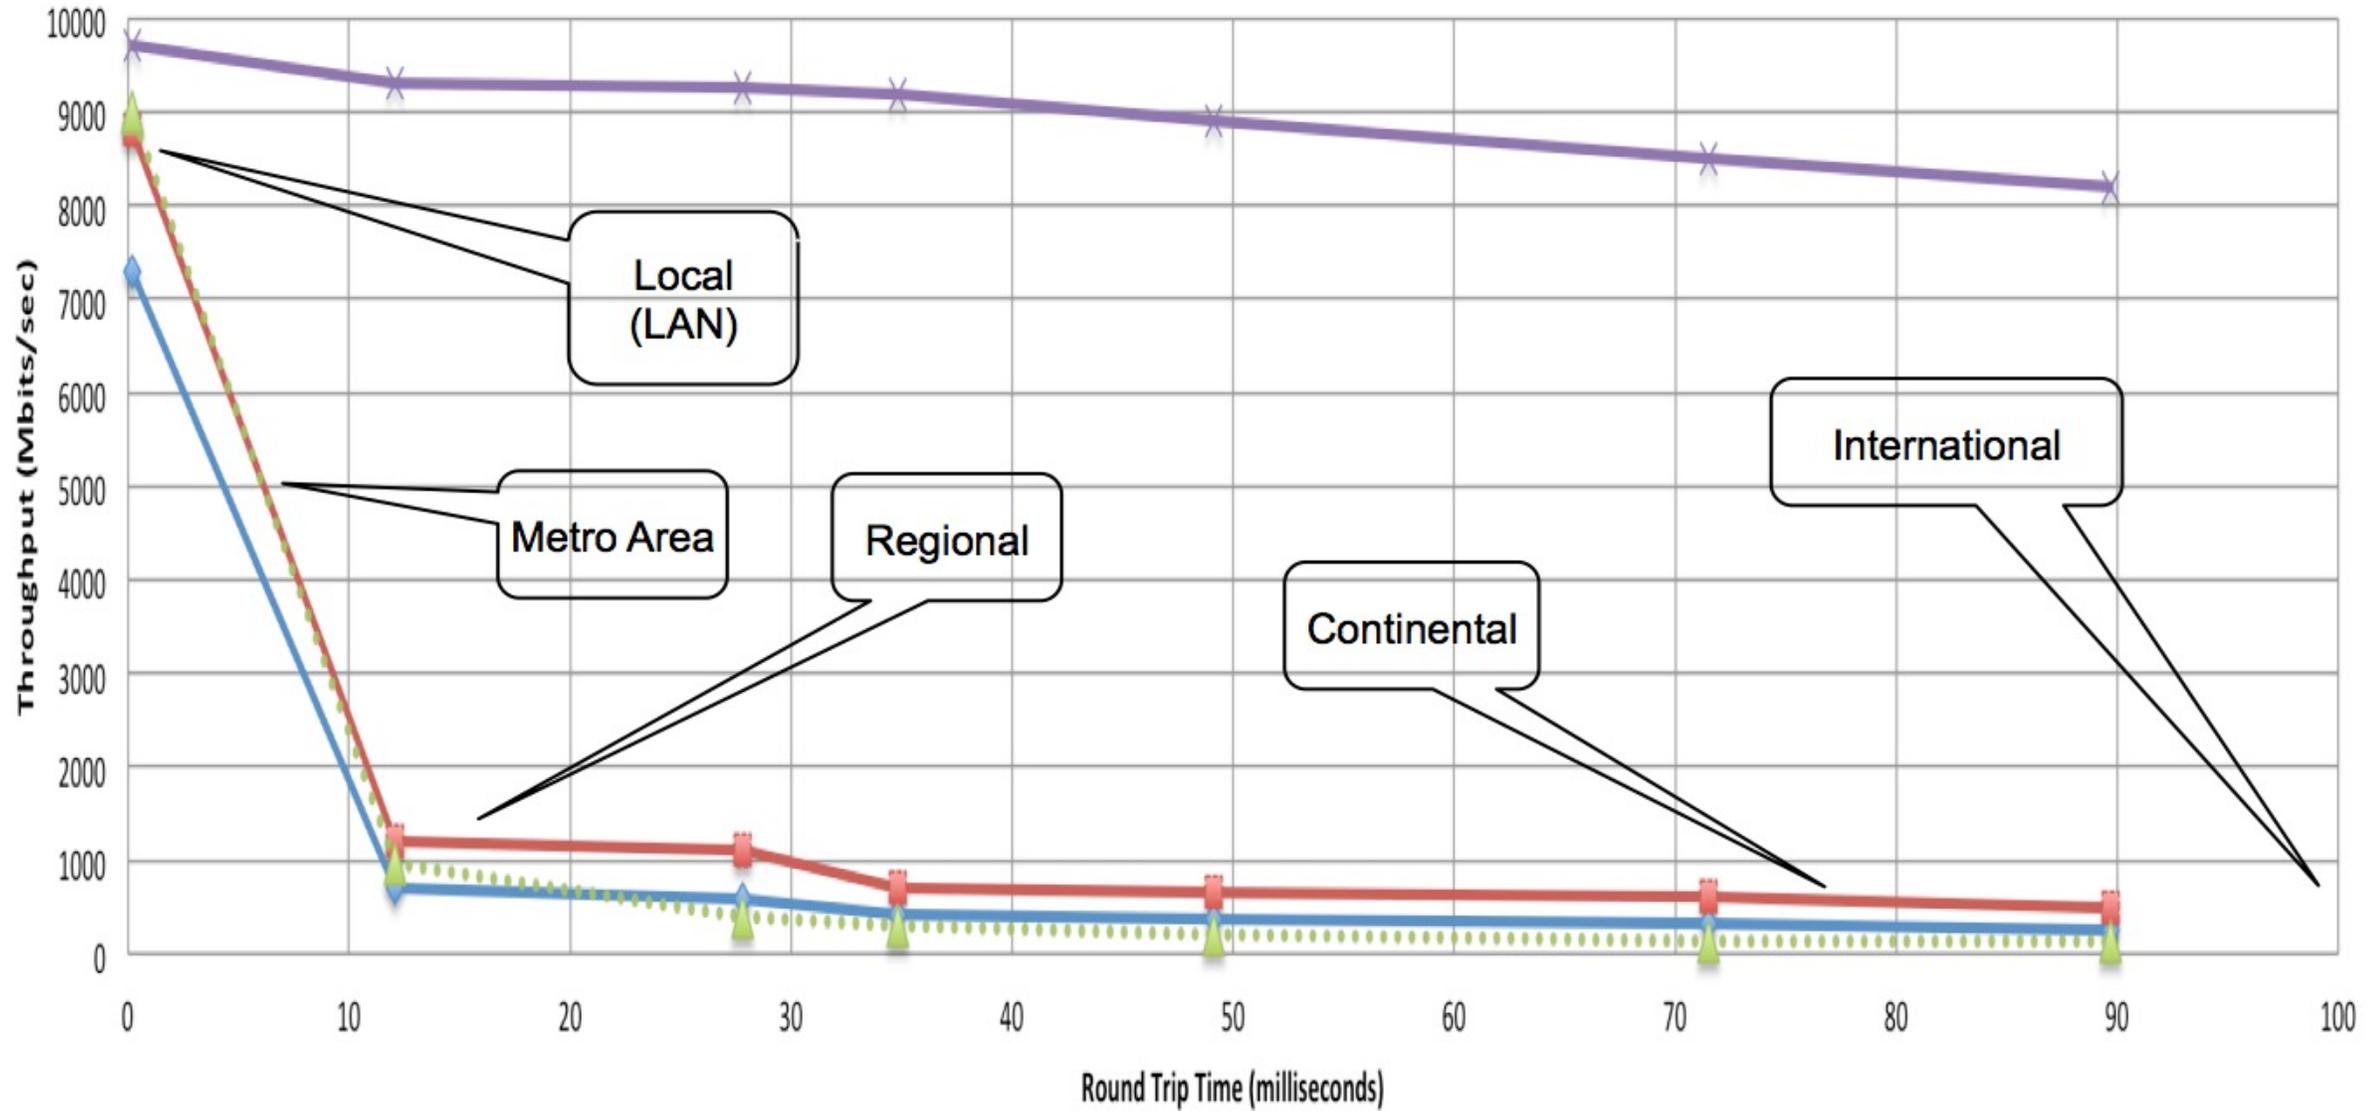
- Latency (or RTT -Round Trip Time)
- Buffer/window size.
- Packet loss

All three are interrelated.





Throughput vs. increasing latency on a 10Gb/s link with 0.0046% packet loss



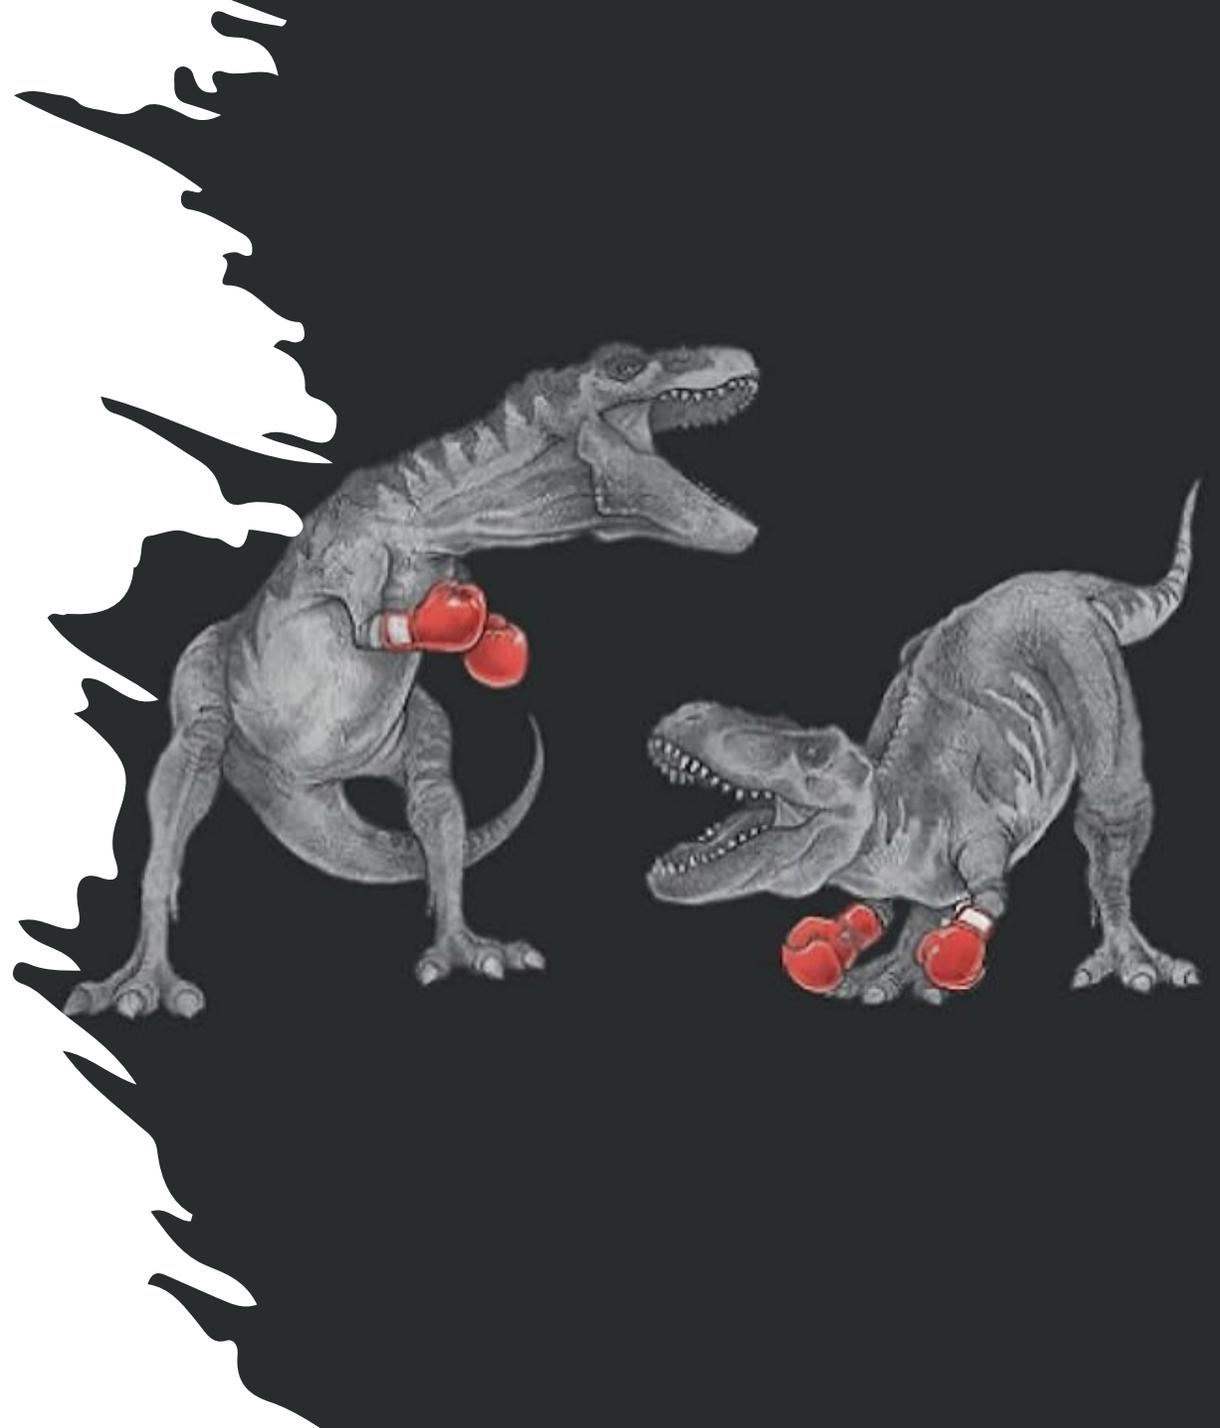
Measured (TCP Reno)

Measured (HTCP)

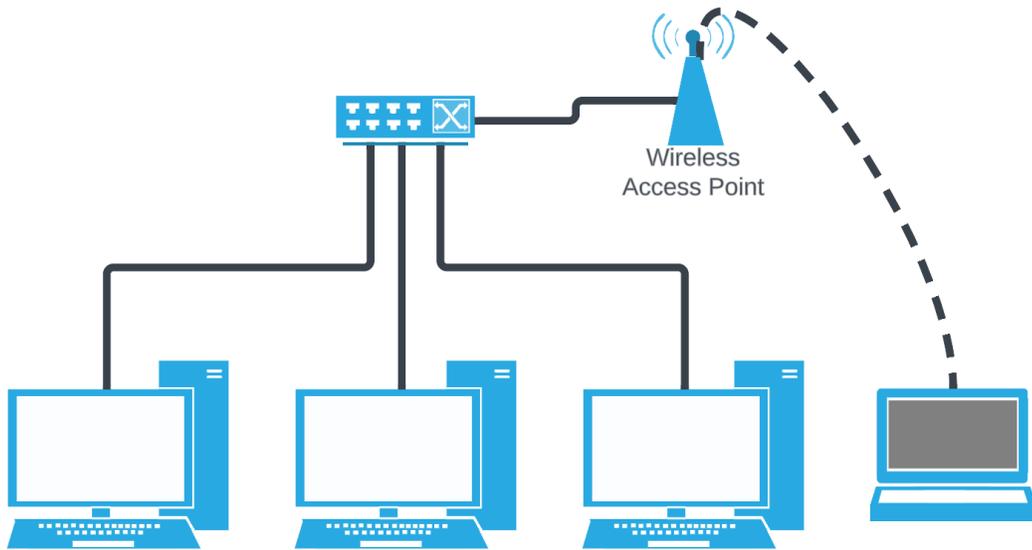
Theoretical (TCP Reno)

Measured (no loss)

WAN vs LAN vs MAN



LAN vs MAN vs WAN



- LAN

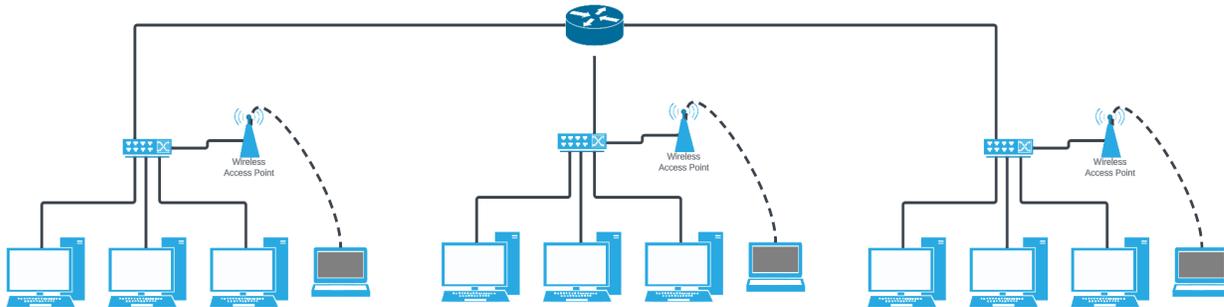
- Local Area Network
- Shorter distances *
- "internal" network *
- Switching and routing *
- You provide *

* As with all networking the "it depends" rule applies

LAN vs MAN vs WAN

- LAN

- Local Area Network
- Shorter distances *
- "internal" network *
- Switching and routing *
- You provide *

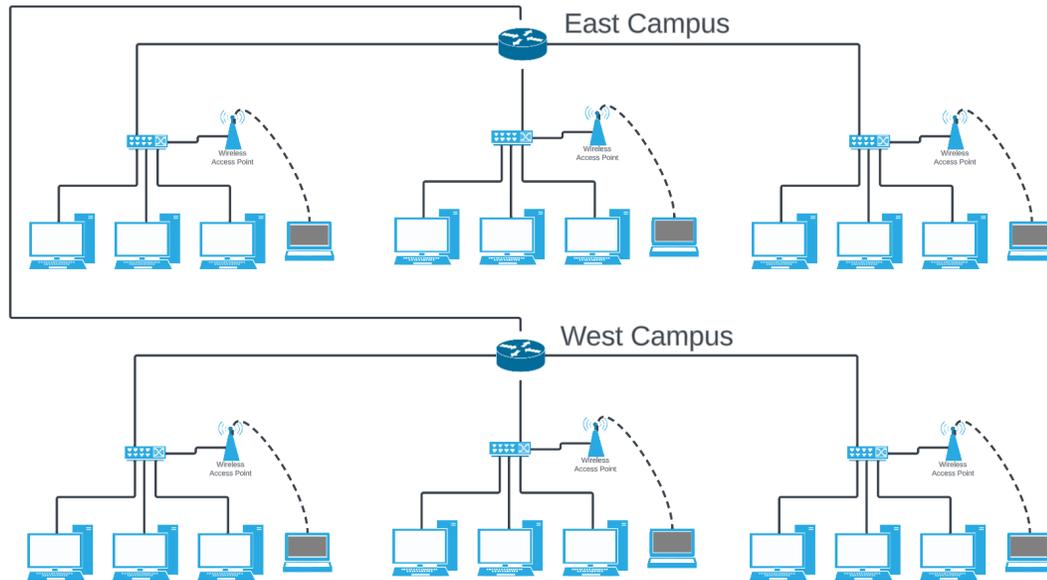


* As with all networking the "it depends" rule applies

LAN vs MAN vs WAN

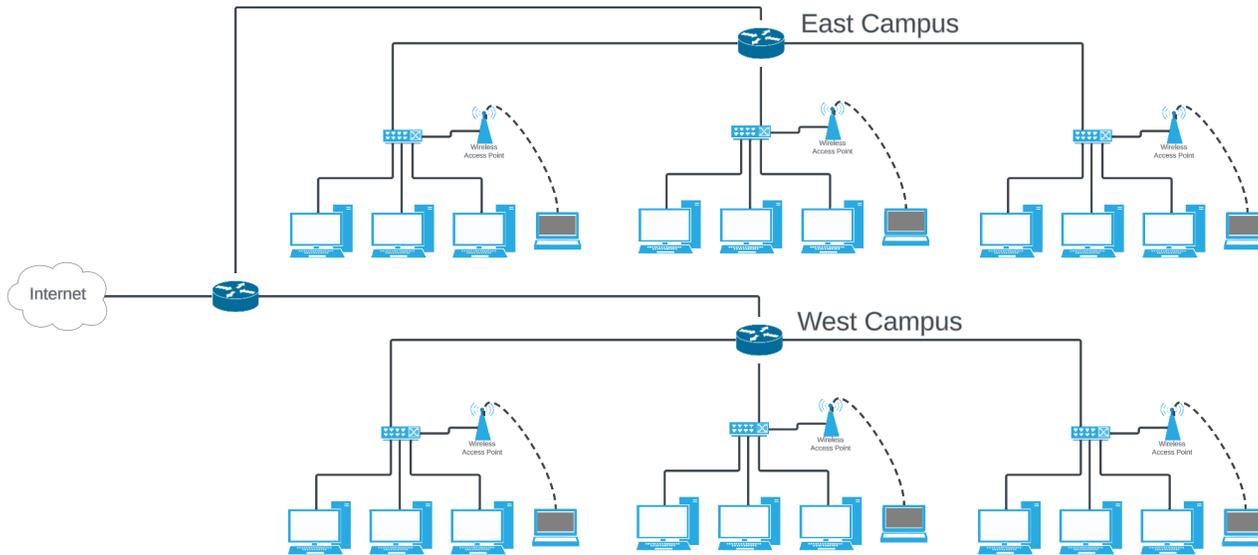
- MAN

- Metropolitan Area Network
- Connections that are between LAN and WAN
- University campuses *



* As with all networking the “it depends” rule applies

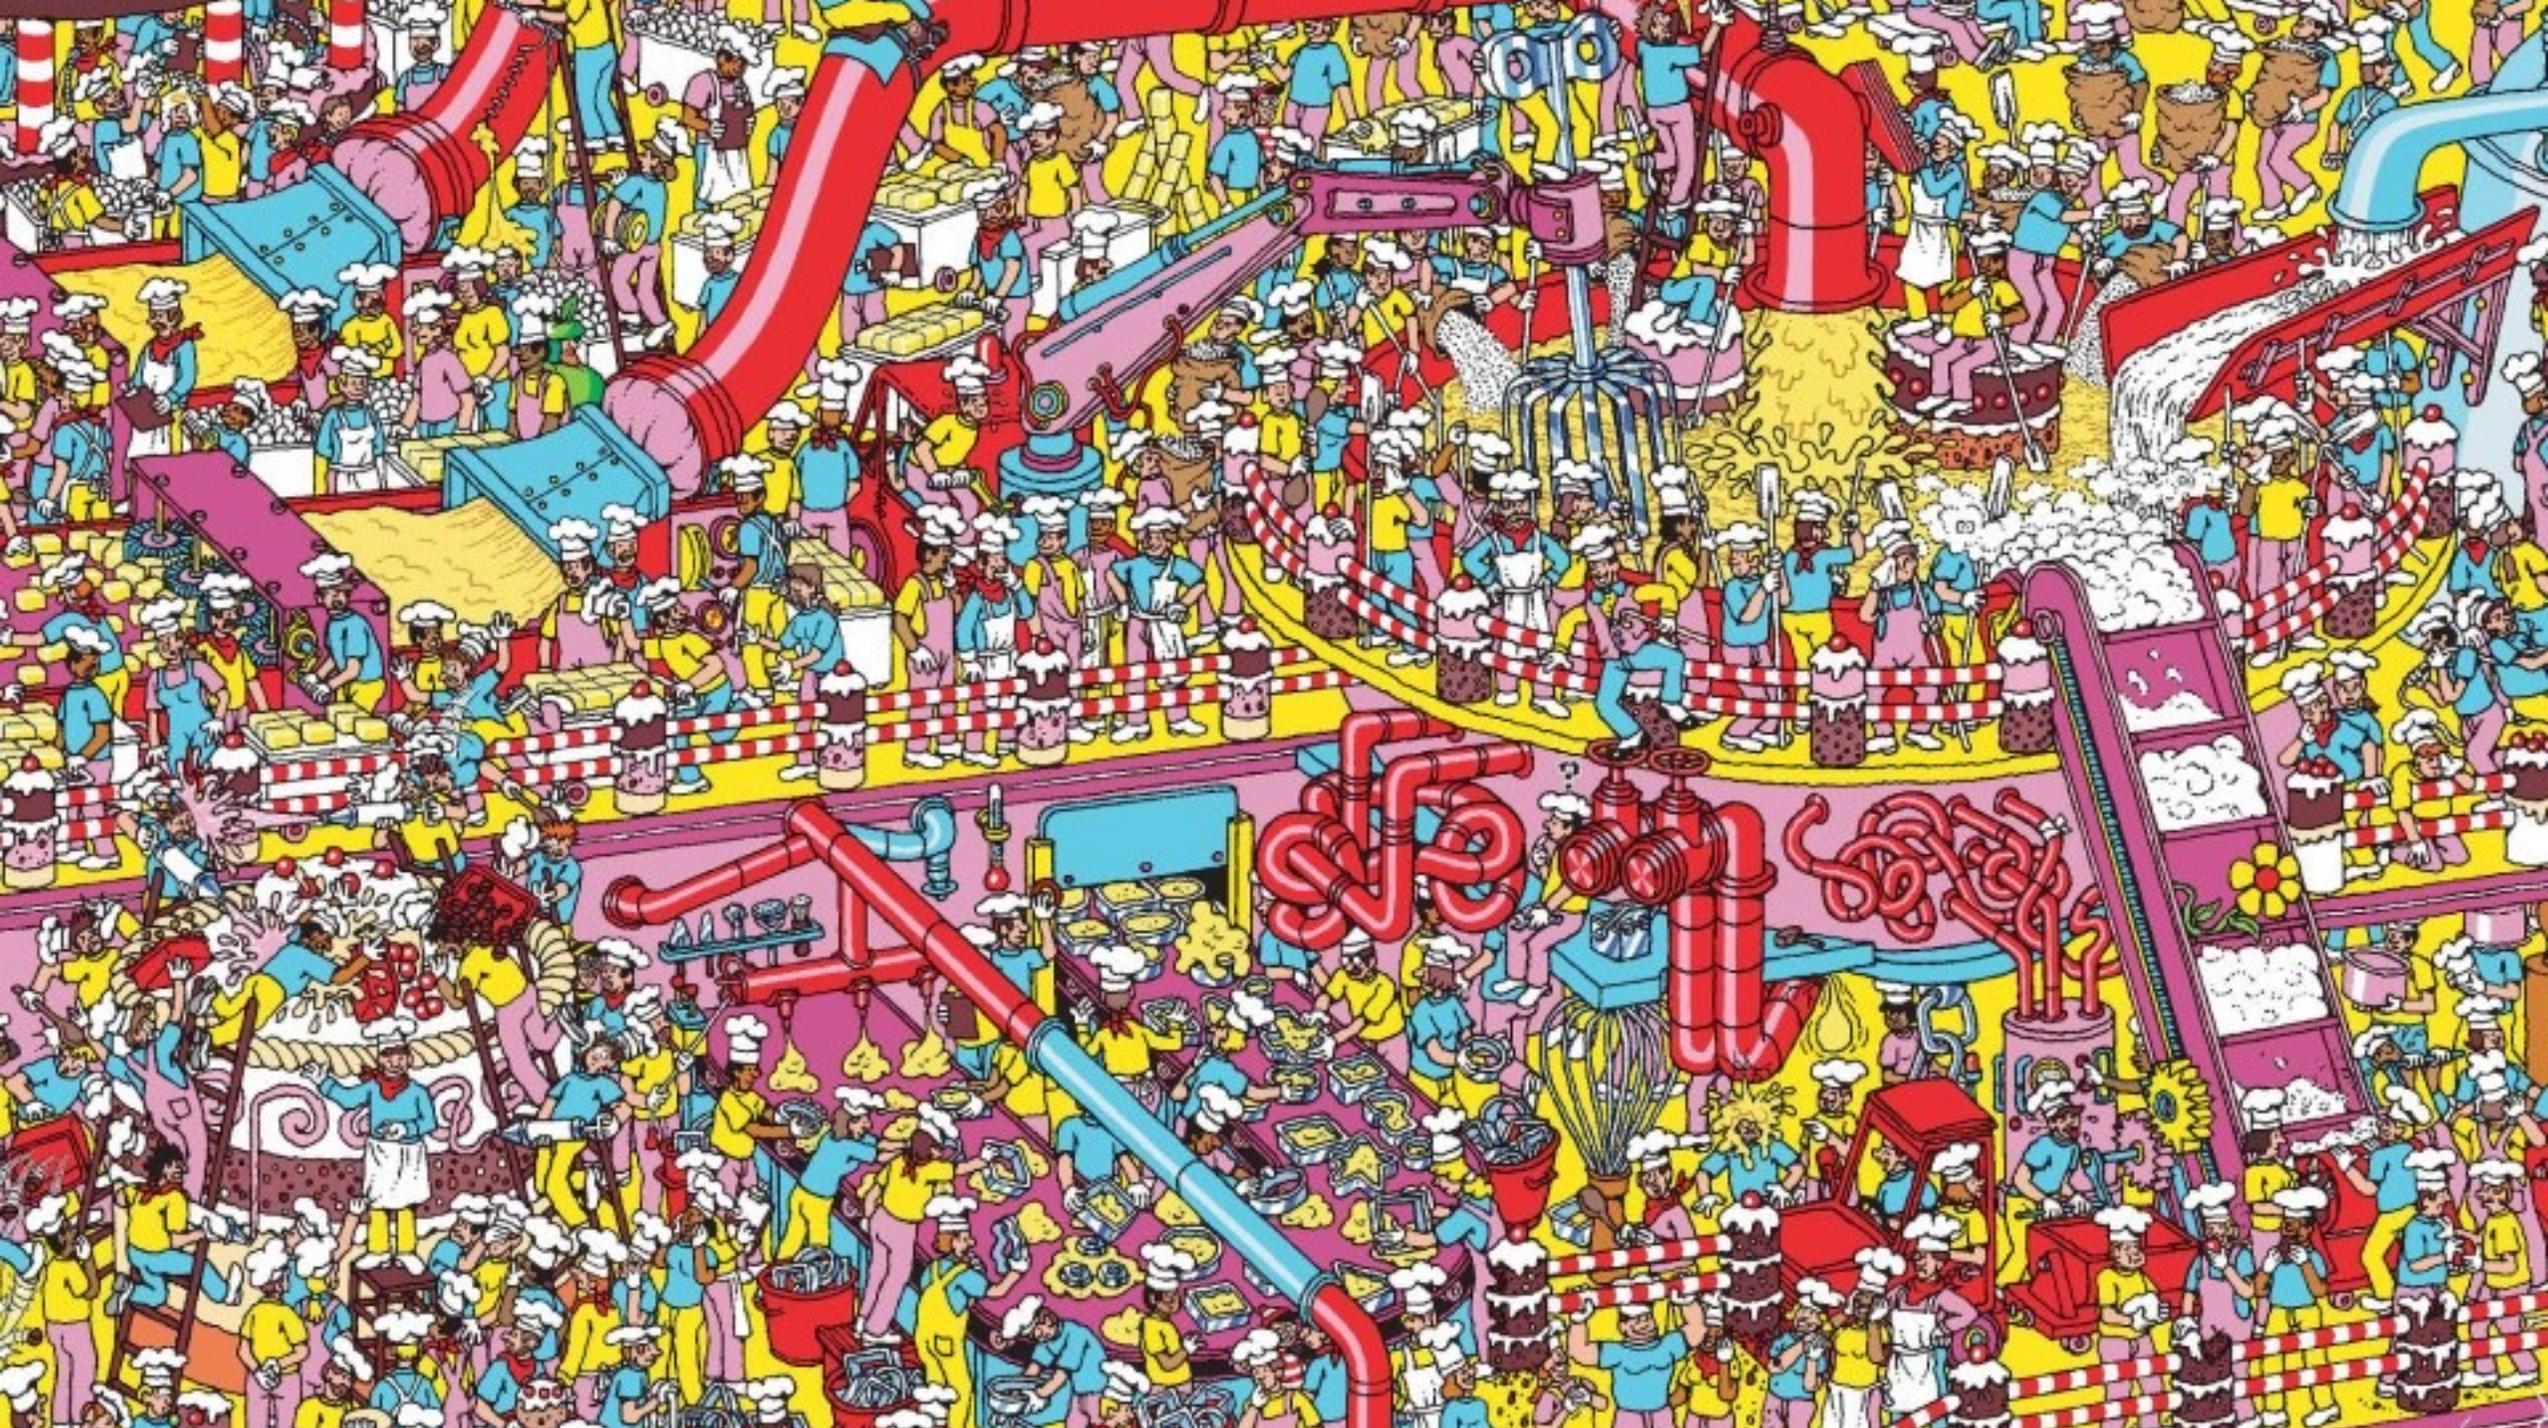
LAN vs MAN vs WAN



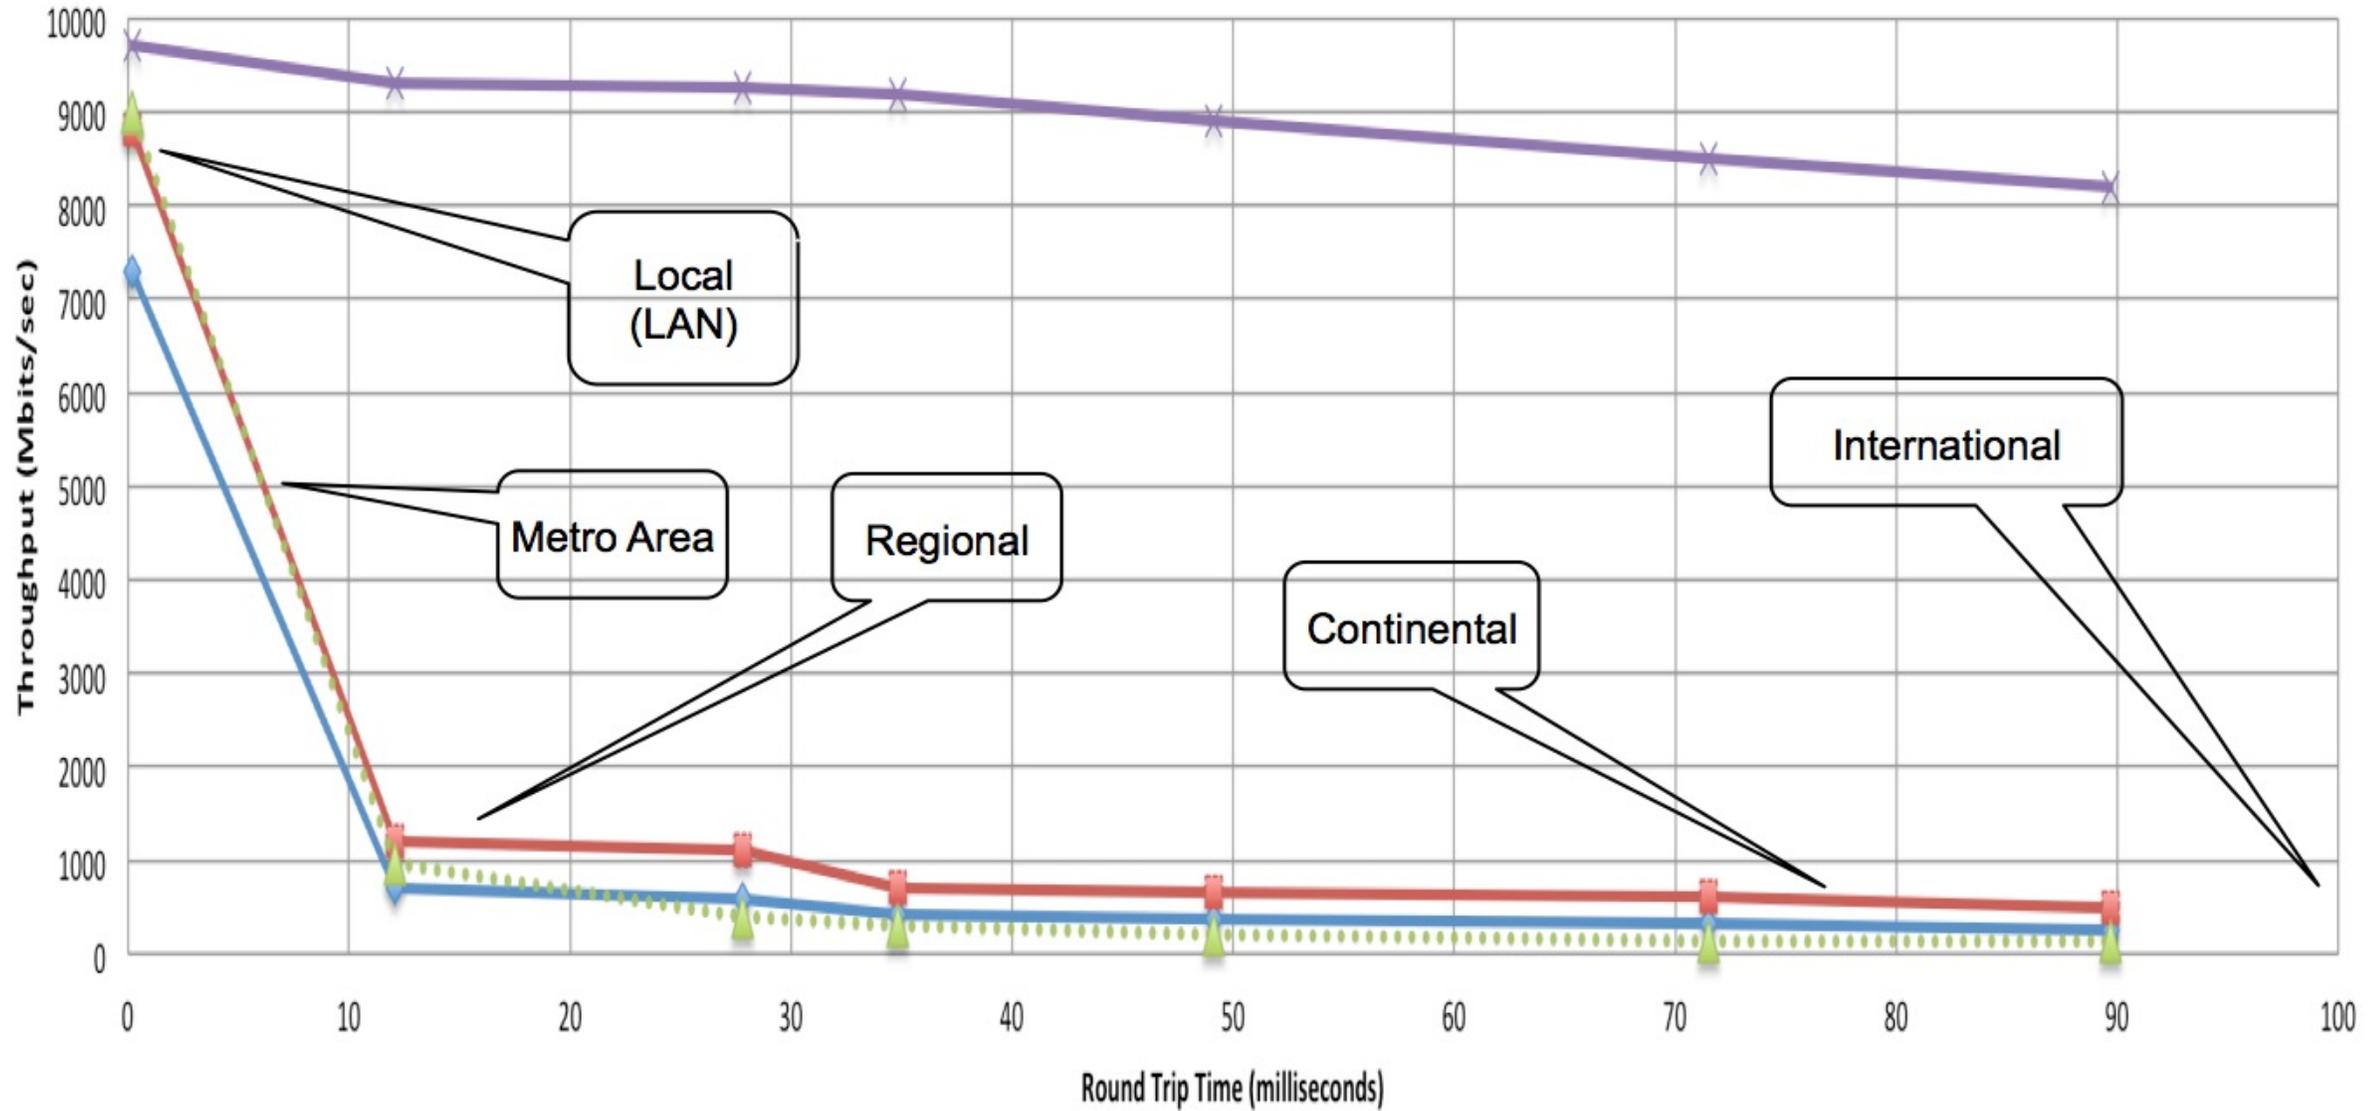
- WAN

- Wide Area Network
- Crosses longer distances *
- Sometimes considered "outside" the network *
- Mostly routing *
- Someone else provides *

* As with all networking the "it depends" rule applies



Throughput vs. increasing latency on a 10Gb/s link with 0.0046% packet loss



Measured (TCP Reno)

Measured (HTCP)

Theoretical (TCP Reno)

Measured (no loss)



So, what happens when research traffic hits normal networks?

16 03 11 Fri 16:45:35

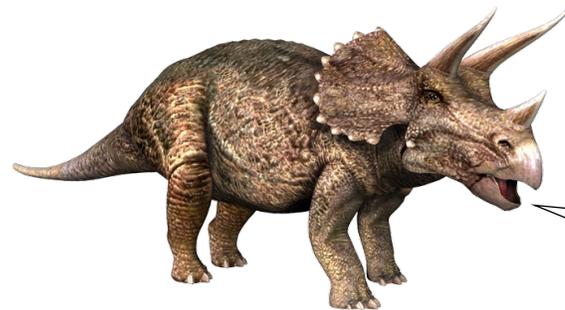


[gifs.com](https://www.gifs.com)

When Data Intensive Science Meets Commercial Commodity Networks

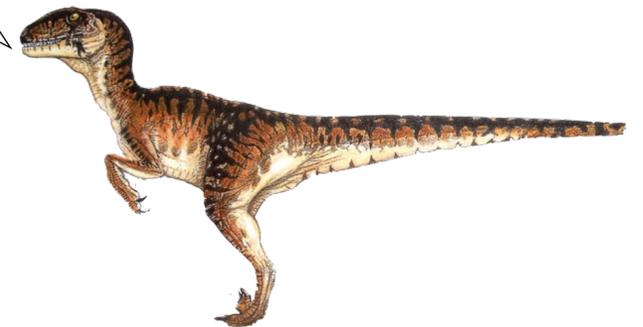
This can result in adverse consequences:

- Performance issues for the researcher
- Performance issues for everyone else
- Frustration for the researcher
- Frustration everyone else
- Frustration for the IT staff



Not a fan of
being lunch...

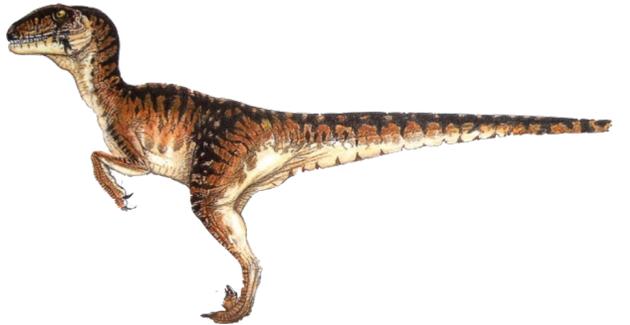
So hungry...



When Data Intensive Science Meets Commercial Commodity Networks

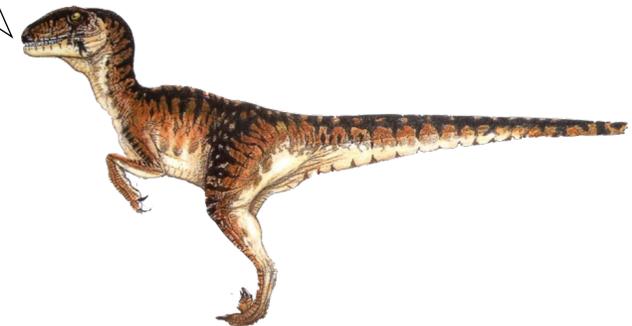


Sigh. I guess cancer cures can wait.



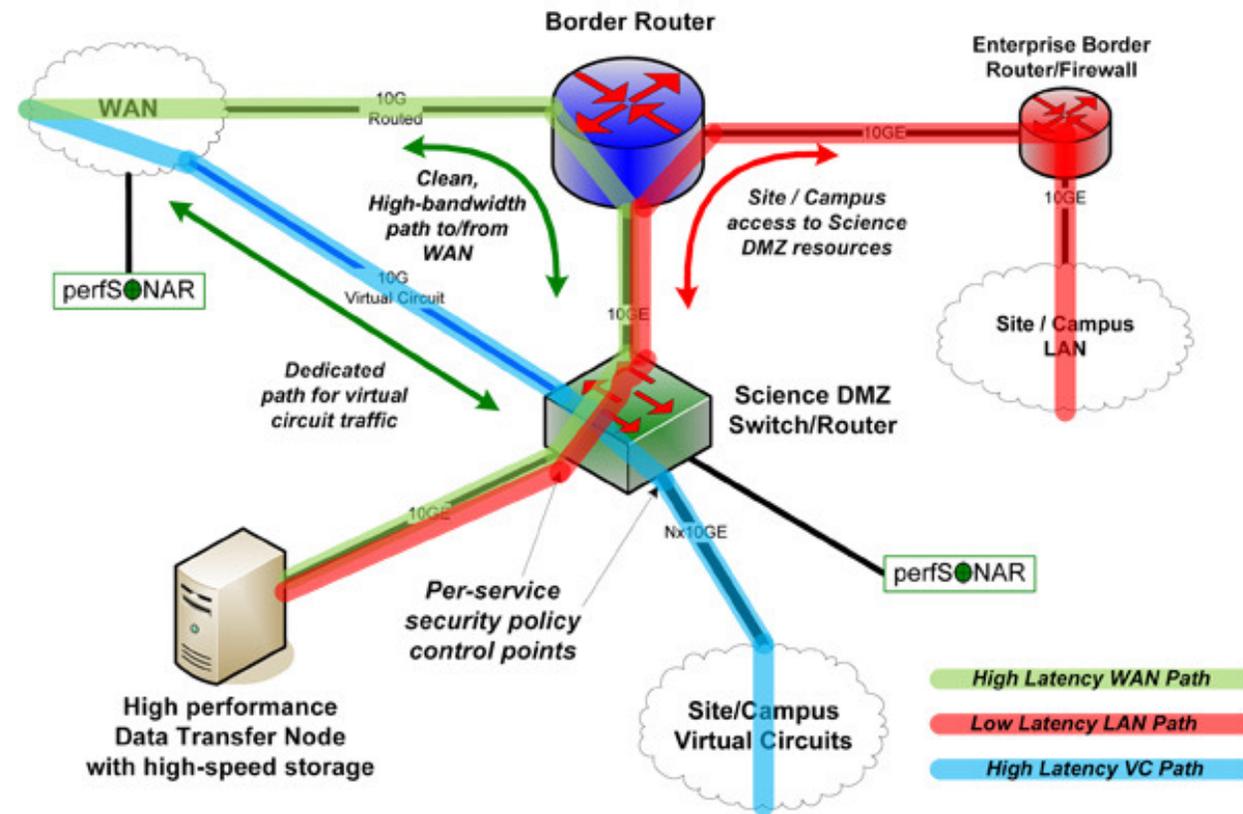
When Data Intensive Science Meets Commercial Commodity Networks

But how do we overcome this? *I can't stop my research* just because **the network can't keep up!** Being able to collaborate is the future of science!



Specialty networks to the rescue!

- Both internally to your organization and externally
- Science DMZ is an example of a specialty network



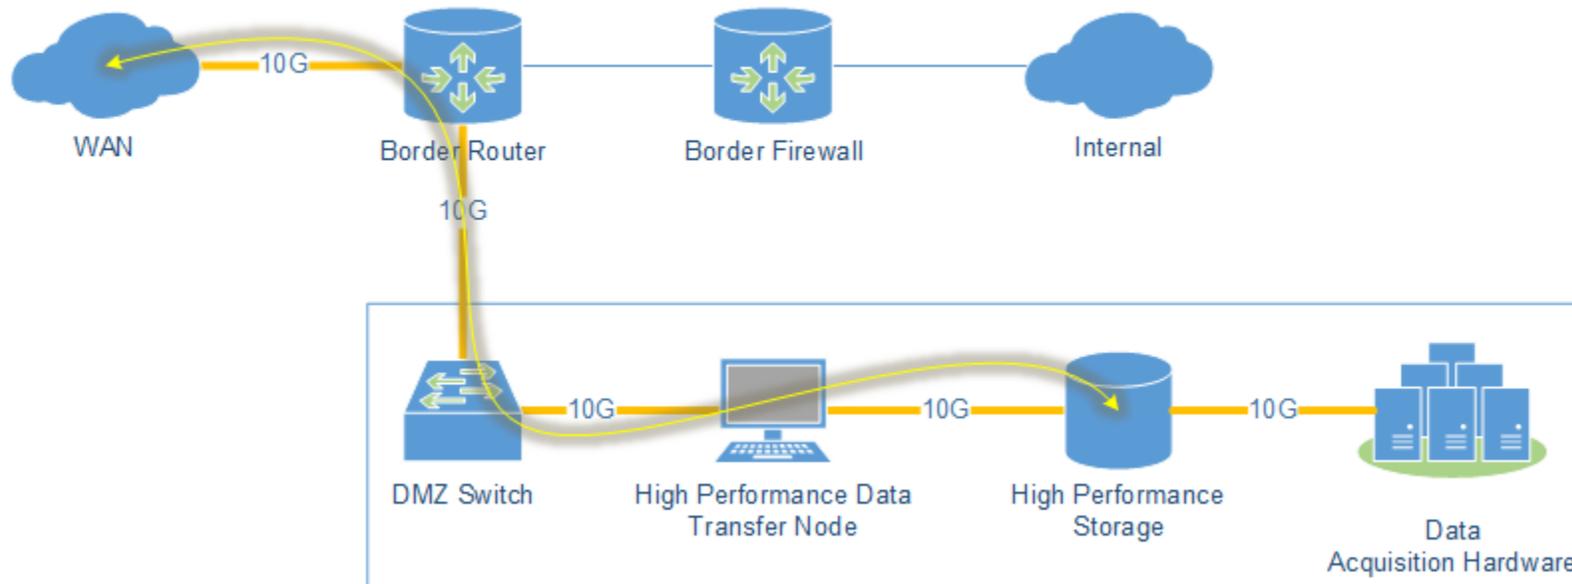
Parts of the puzzle...

- Science DMZ
- DTN (Data Transfer Node)
- Storage
- Research and Education networks
 - NREN (National Research and Education Network)
 - RON (Regional Optical Network)
 - R&E Network (Research and Education Network)



Data Transfer Nodes (DTN)

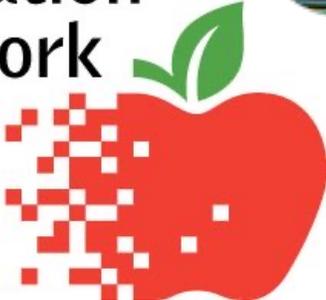
- Simply put, a DTN is a server that is specially designed to move data from disk to a network at speed.



State/regional networks aka "your ISP"

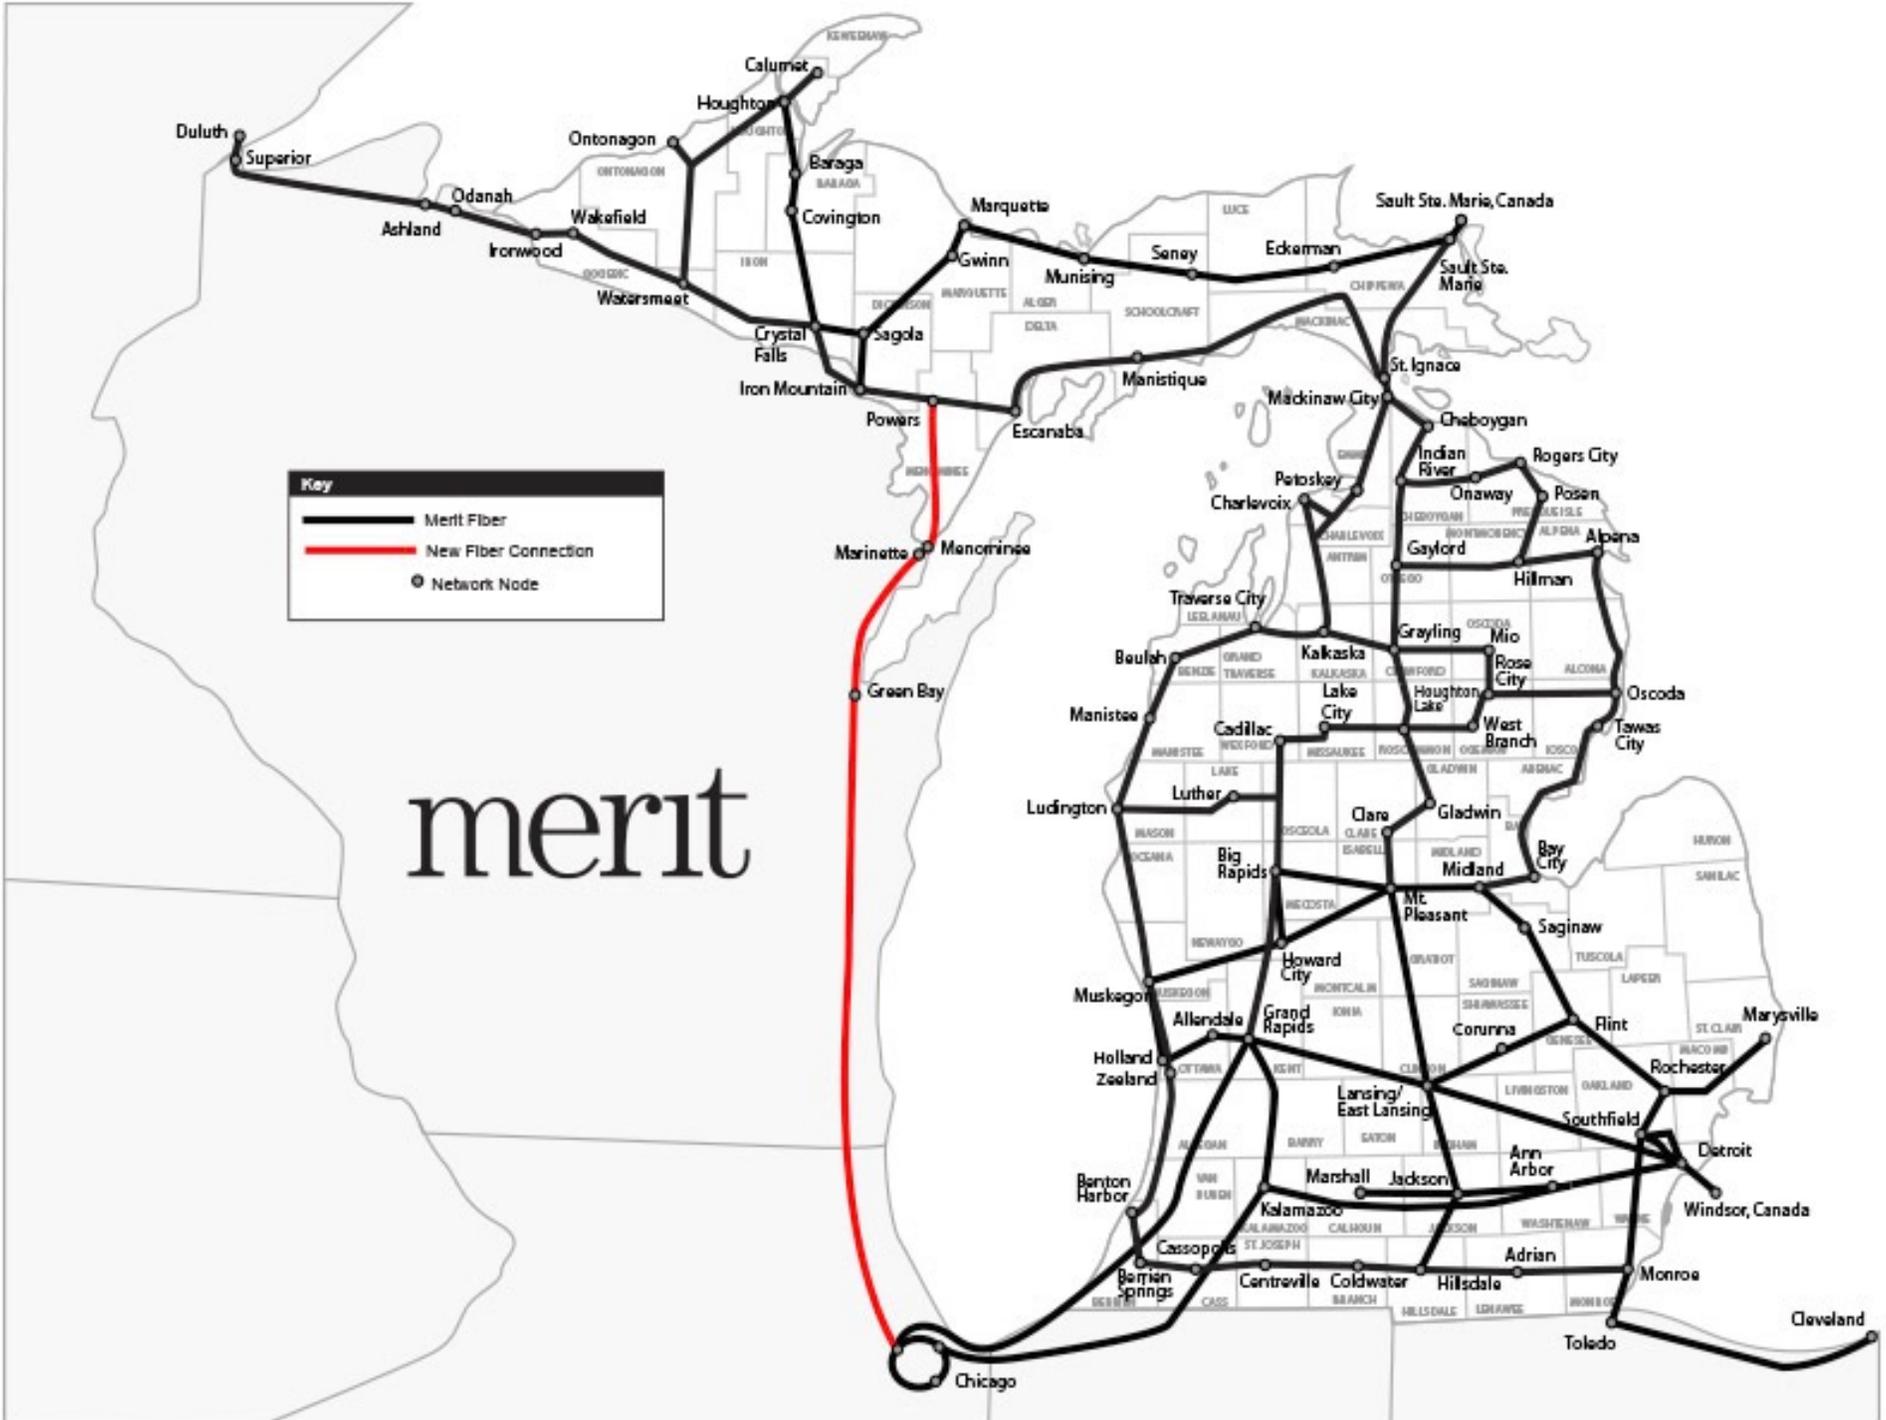


K-20
Education
Network



Relationships make the packets flow...





Key

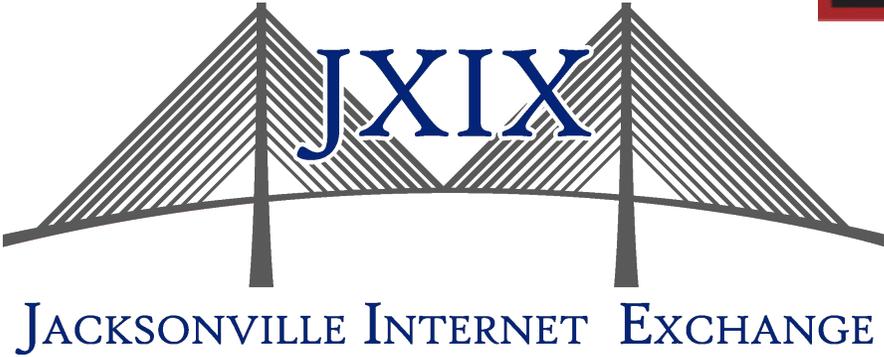
- Merit Fiber
- New Fiber Connection
- Network Node

merit

Exchanges and PoPs



REANNZ



NNENIX
NORTHERN NEW ENGLAND
NEUTRAL INTERNET EXCHANGE

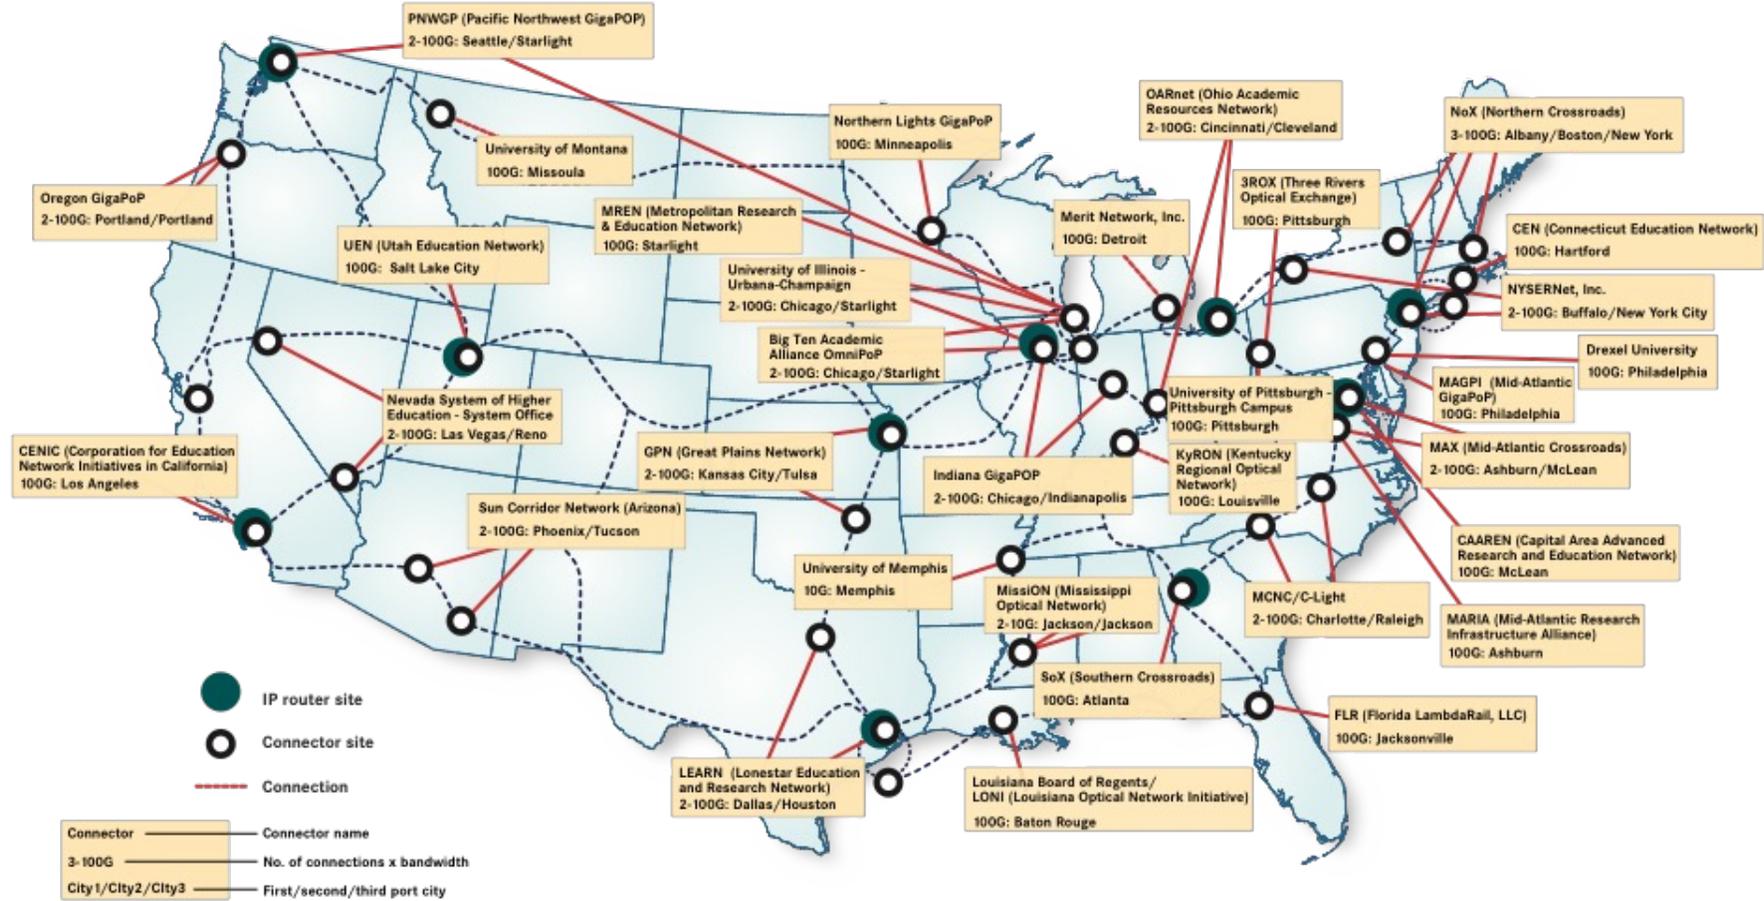
National Networks (US)





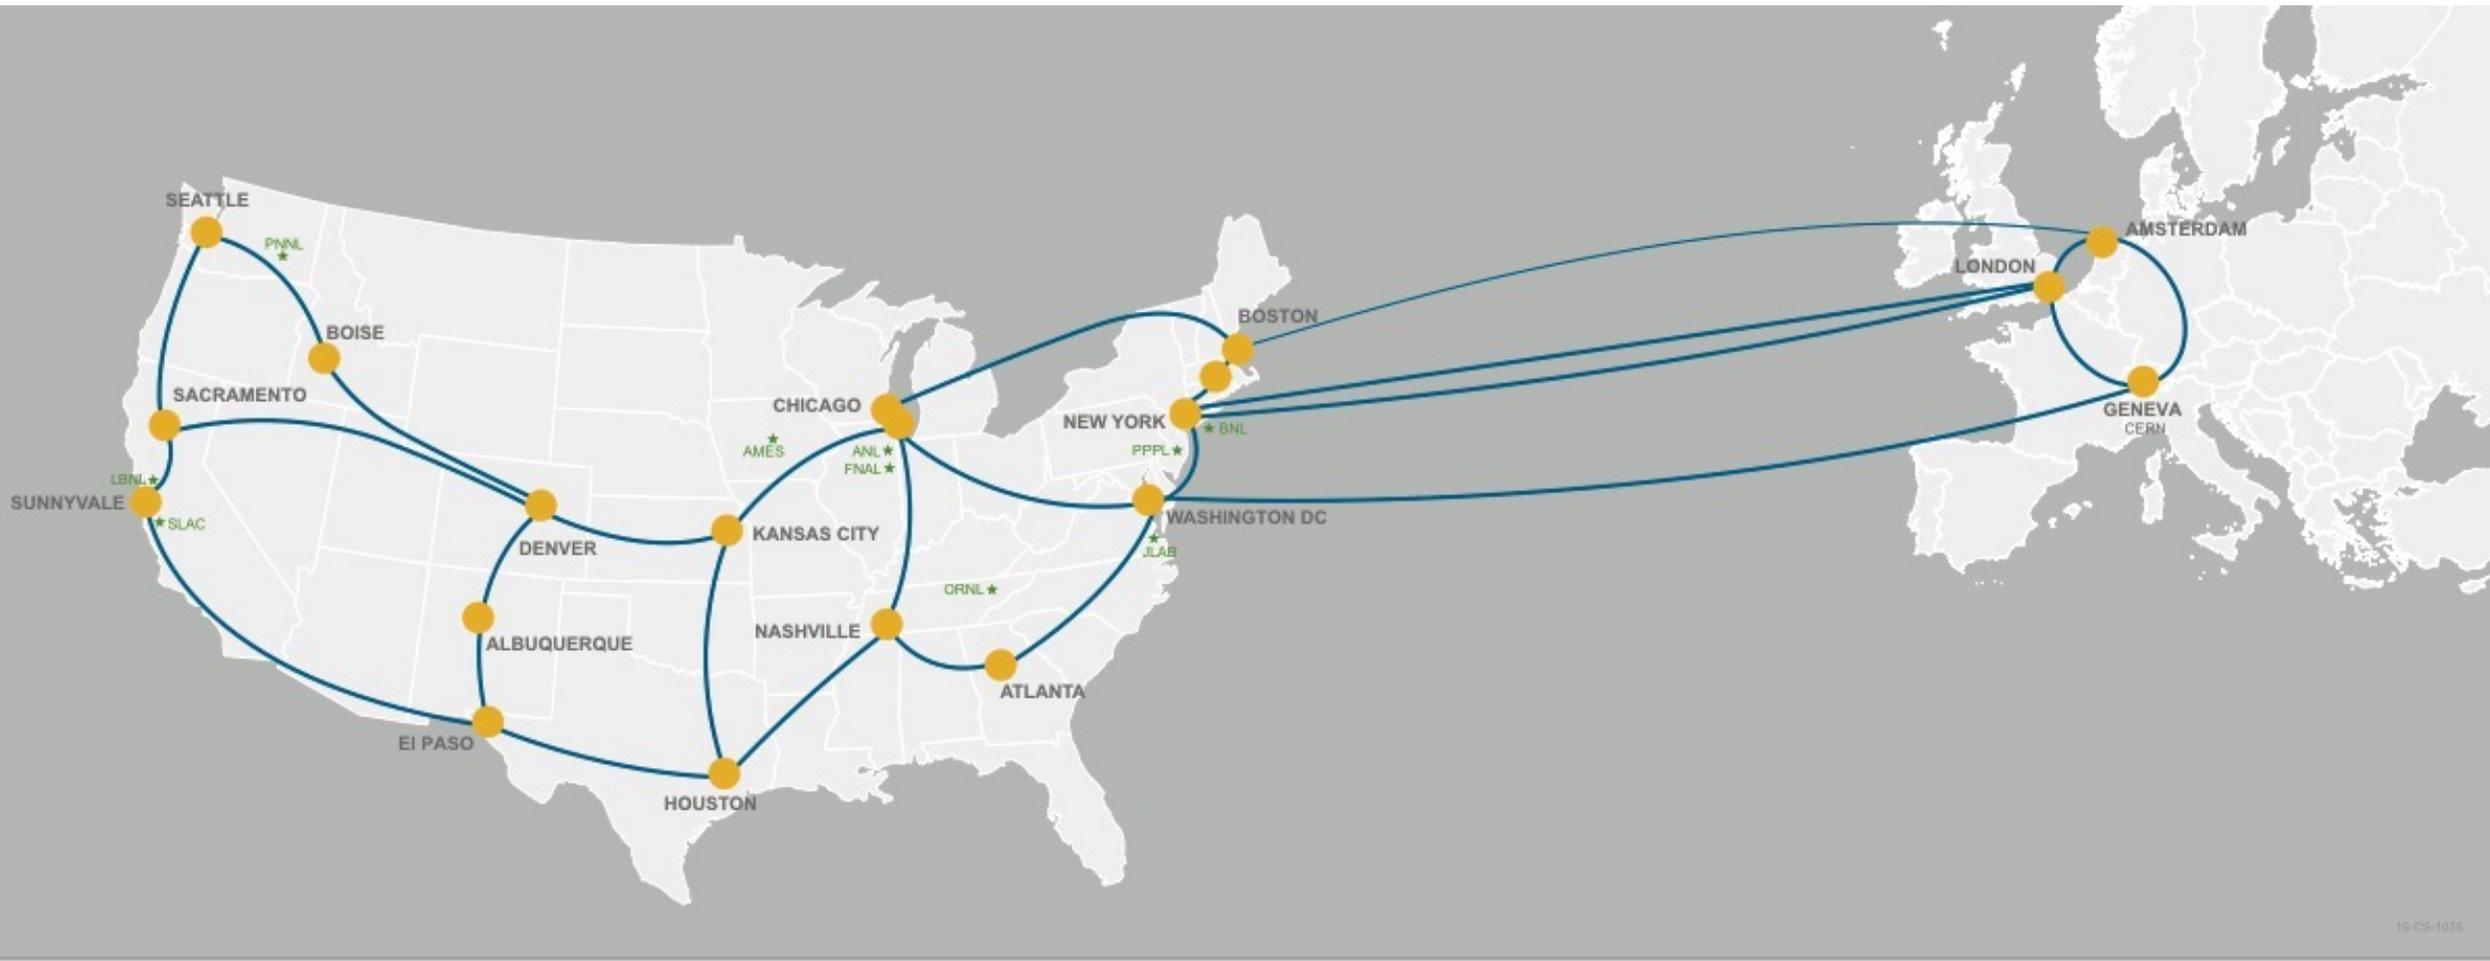
INTERNET2 NETWORK CONNECTIONS

WWW.INTERNET2.EDU/CONNECTORS - MARCH OF 2017



NETWORK PARTNERS





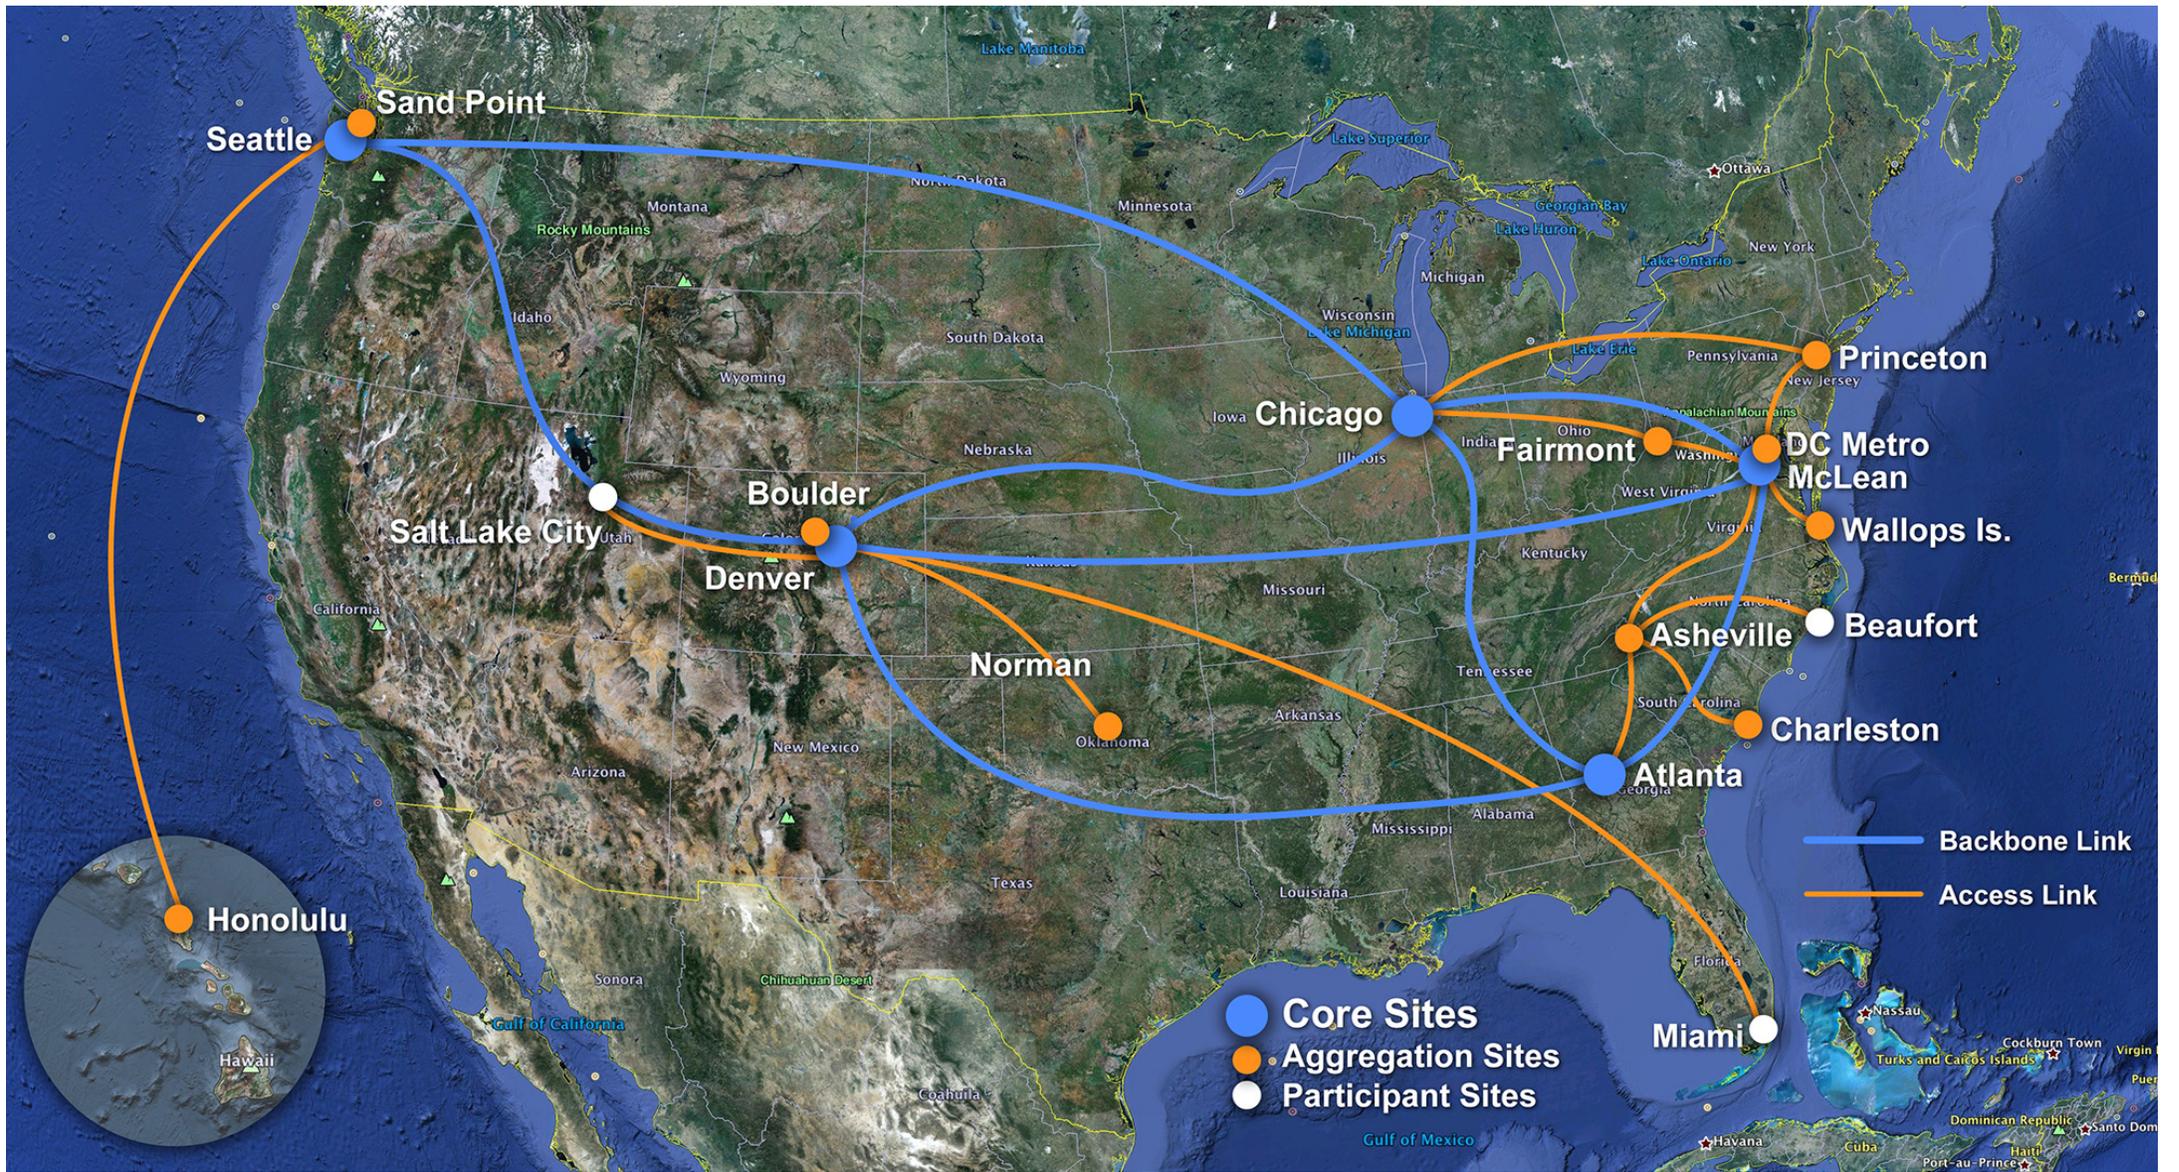
ESnet

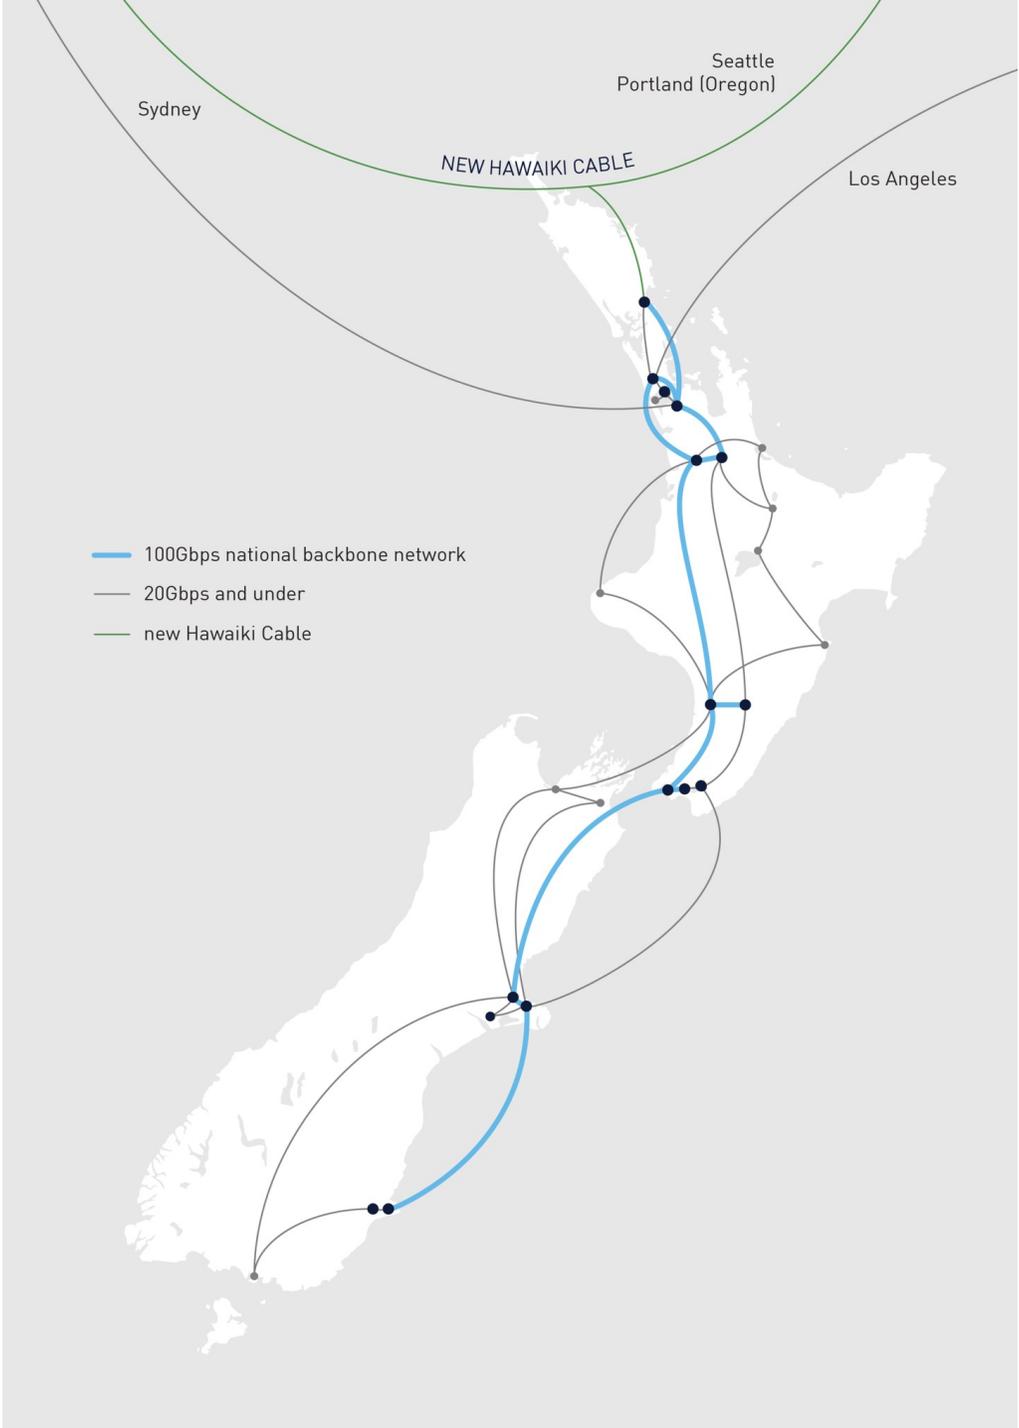
ENERGY SCIENCES NETWORK

★ Department of Energy Office of Science National Labs

- Ames** Ames Laboratory (Ames, IA)
- ANL** Argonne National Laboratory (Argonne, IL)
- BNL** Brookhaven National Laboratory (Upton, NY)
- FNAL** Fermi National Accelerator Laboratory (Batavia, IL)
- JLAB** Thomas Jefferson National Accelerator Facility (Newport News, VA)

- LBL** Lawrence Berkeley National Laboratory (Berkeley, CA)
- ORNL** Oak Ridge National Laboratory (Oak Ridge, TN)
- PNNL** Pacific Northwest National Laboratory (Richland, WA)
- PPPL** Princeton Plasma Physics Laboratory (Princeton, NJ)
- SLAC** SLAC National Accelerator Laboratory (Menlo Park, CA)





- 100Gbps national backbone network
- 20Gbps and under
- new Hawaiki Cable

Sydney

Seattle
Portland (Oregon)

Los Angeles

NEW HAWAIKI CABLE

National Networks

REANVZ



Spanish
Clinical
Research
Network

★
TEIN3

SANReN
South African National
Research Network



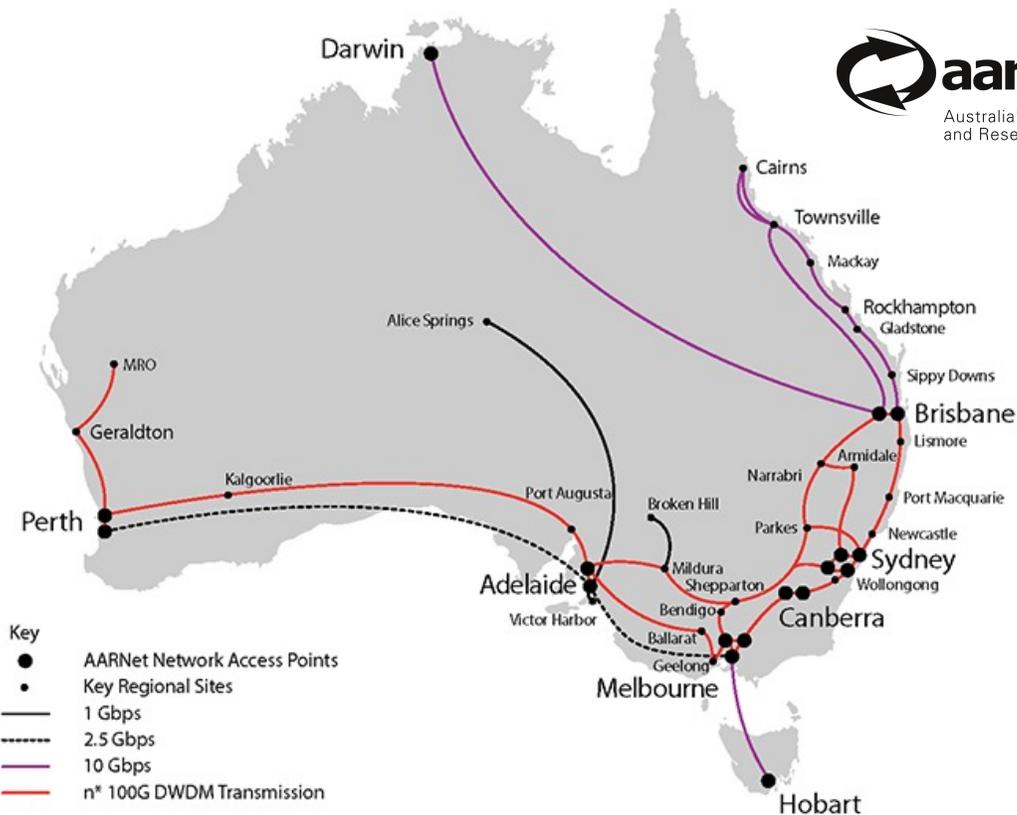
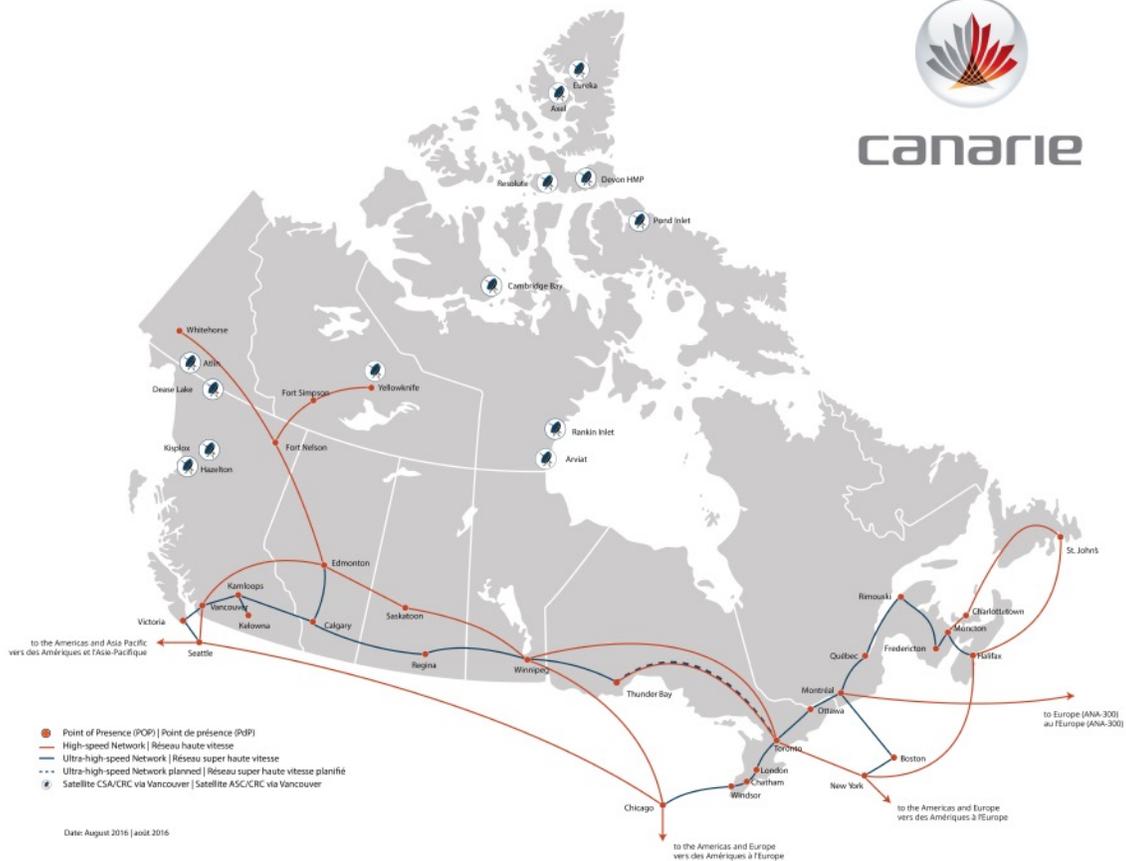
canarie

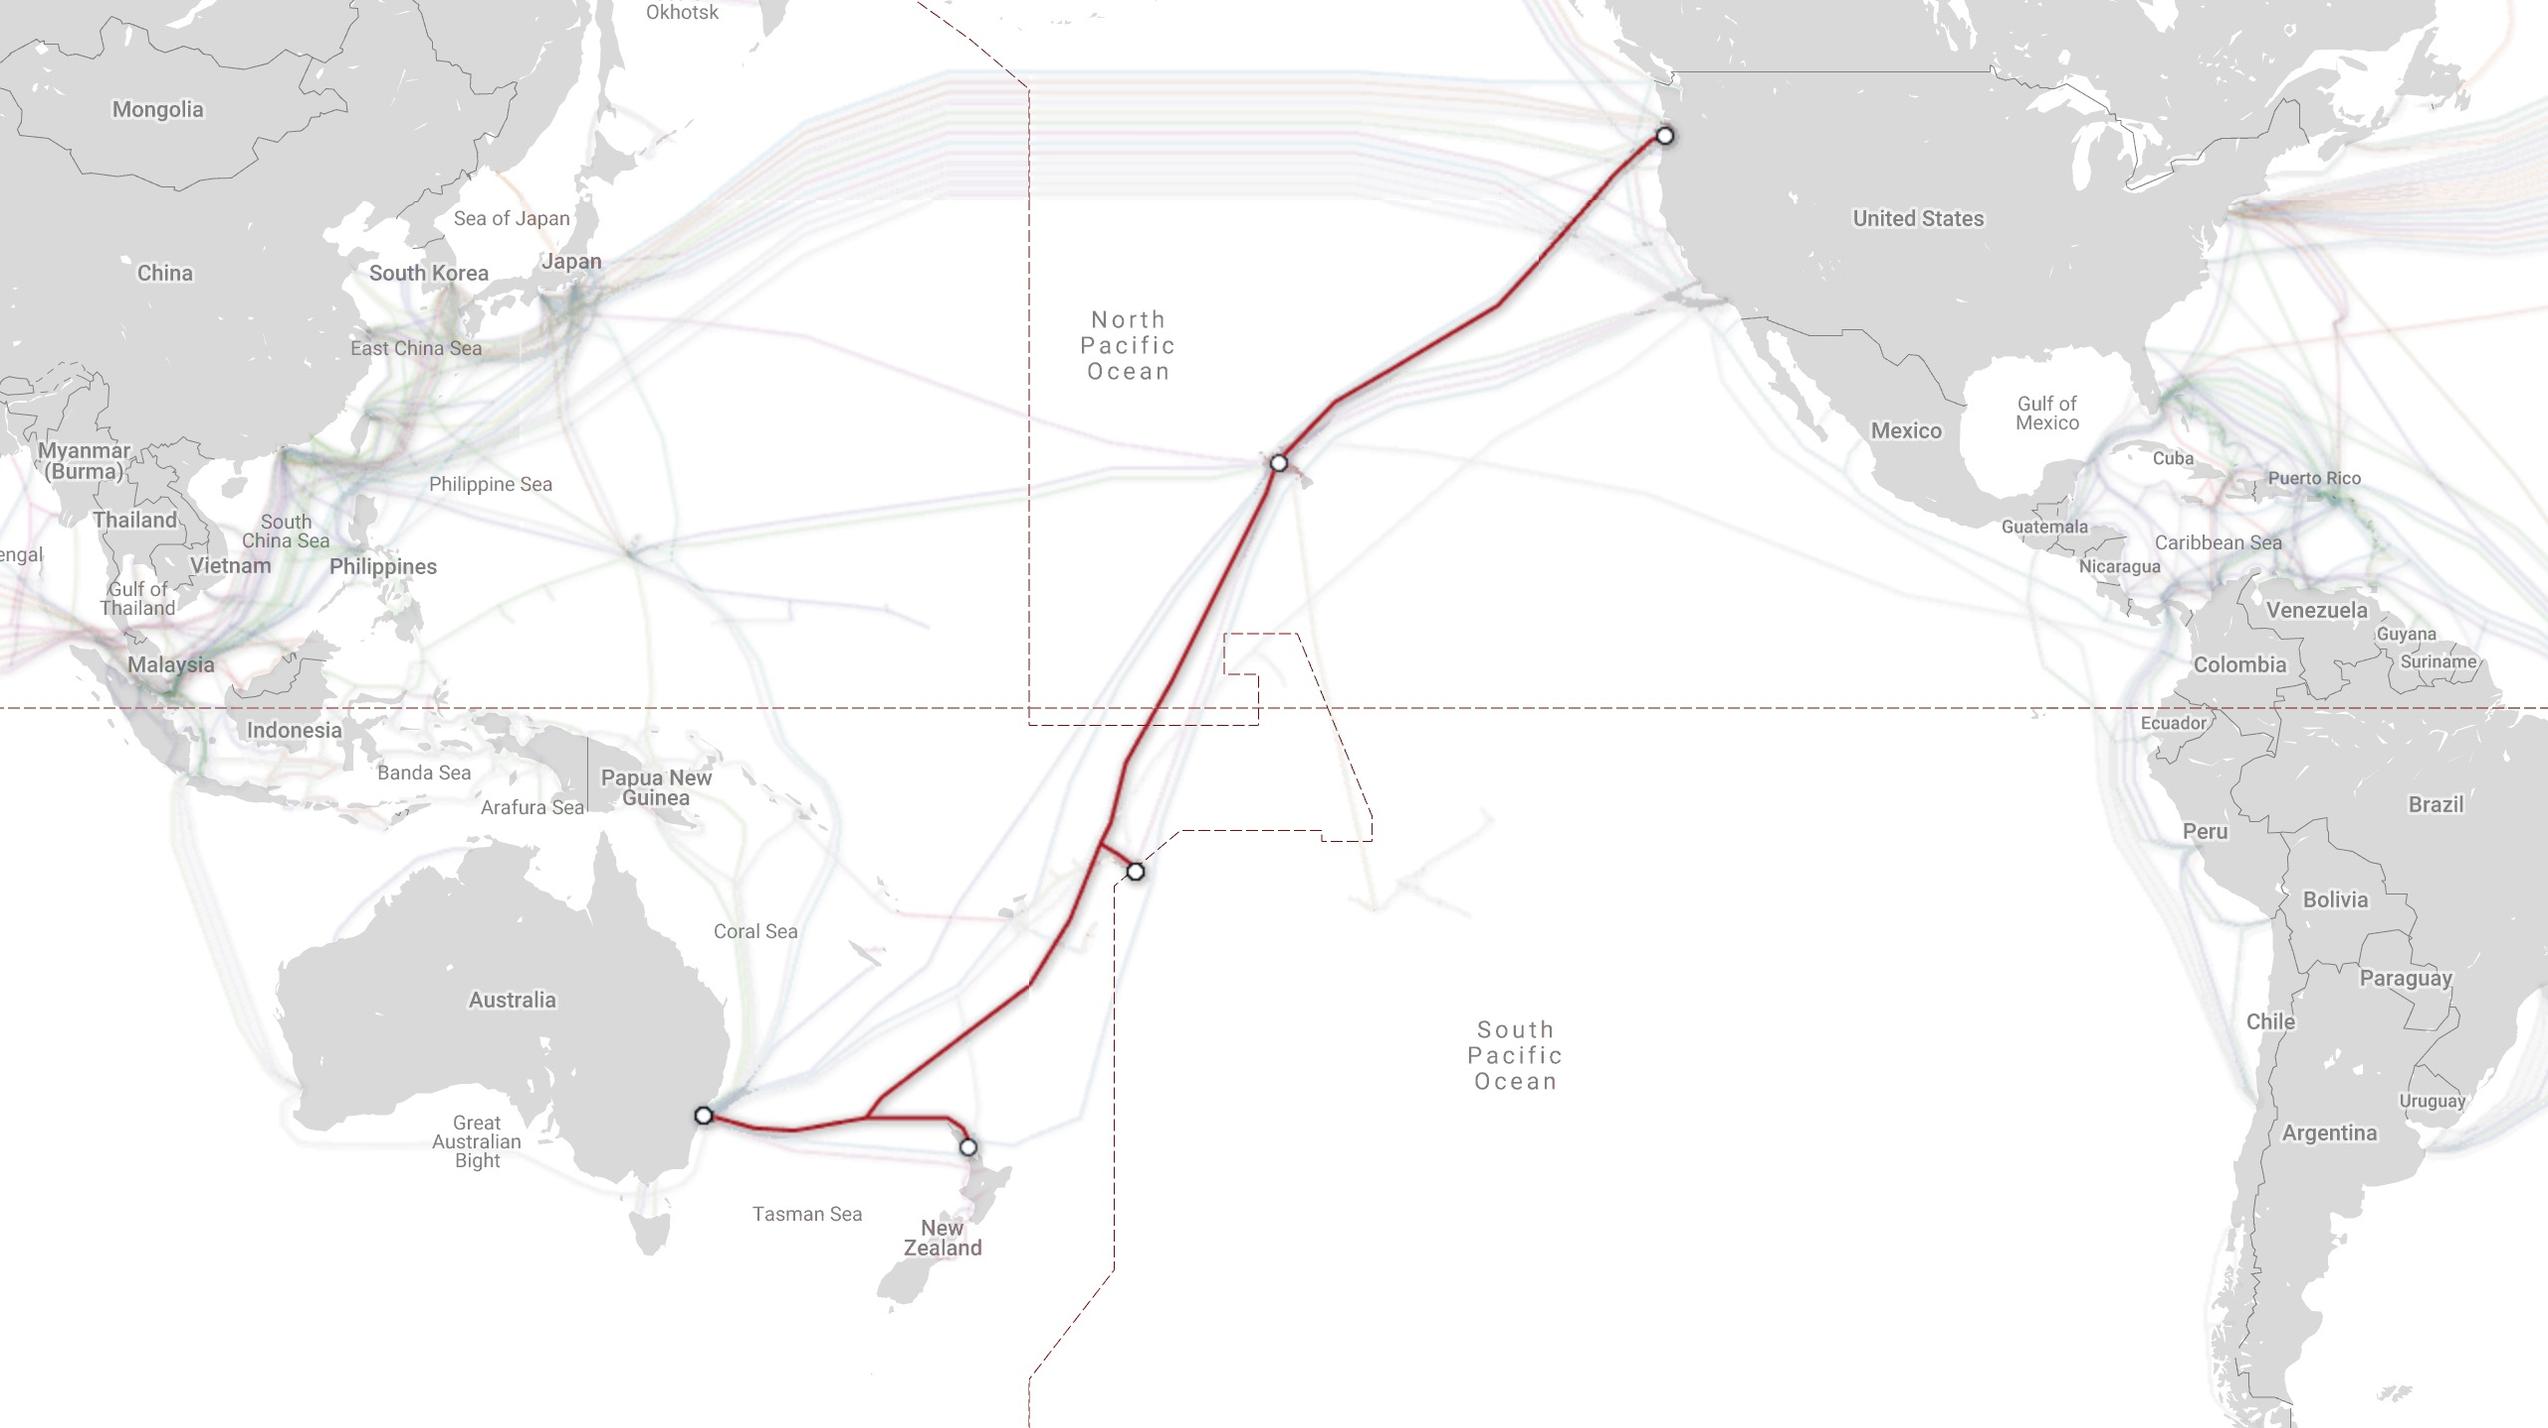


ERNET
Indian Research & Education Network



Australia's Academic
and Research Network



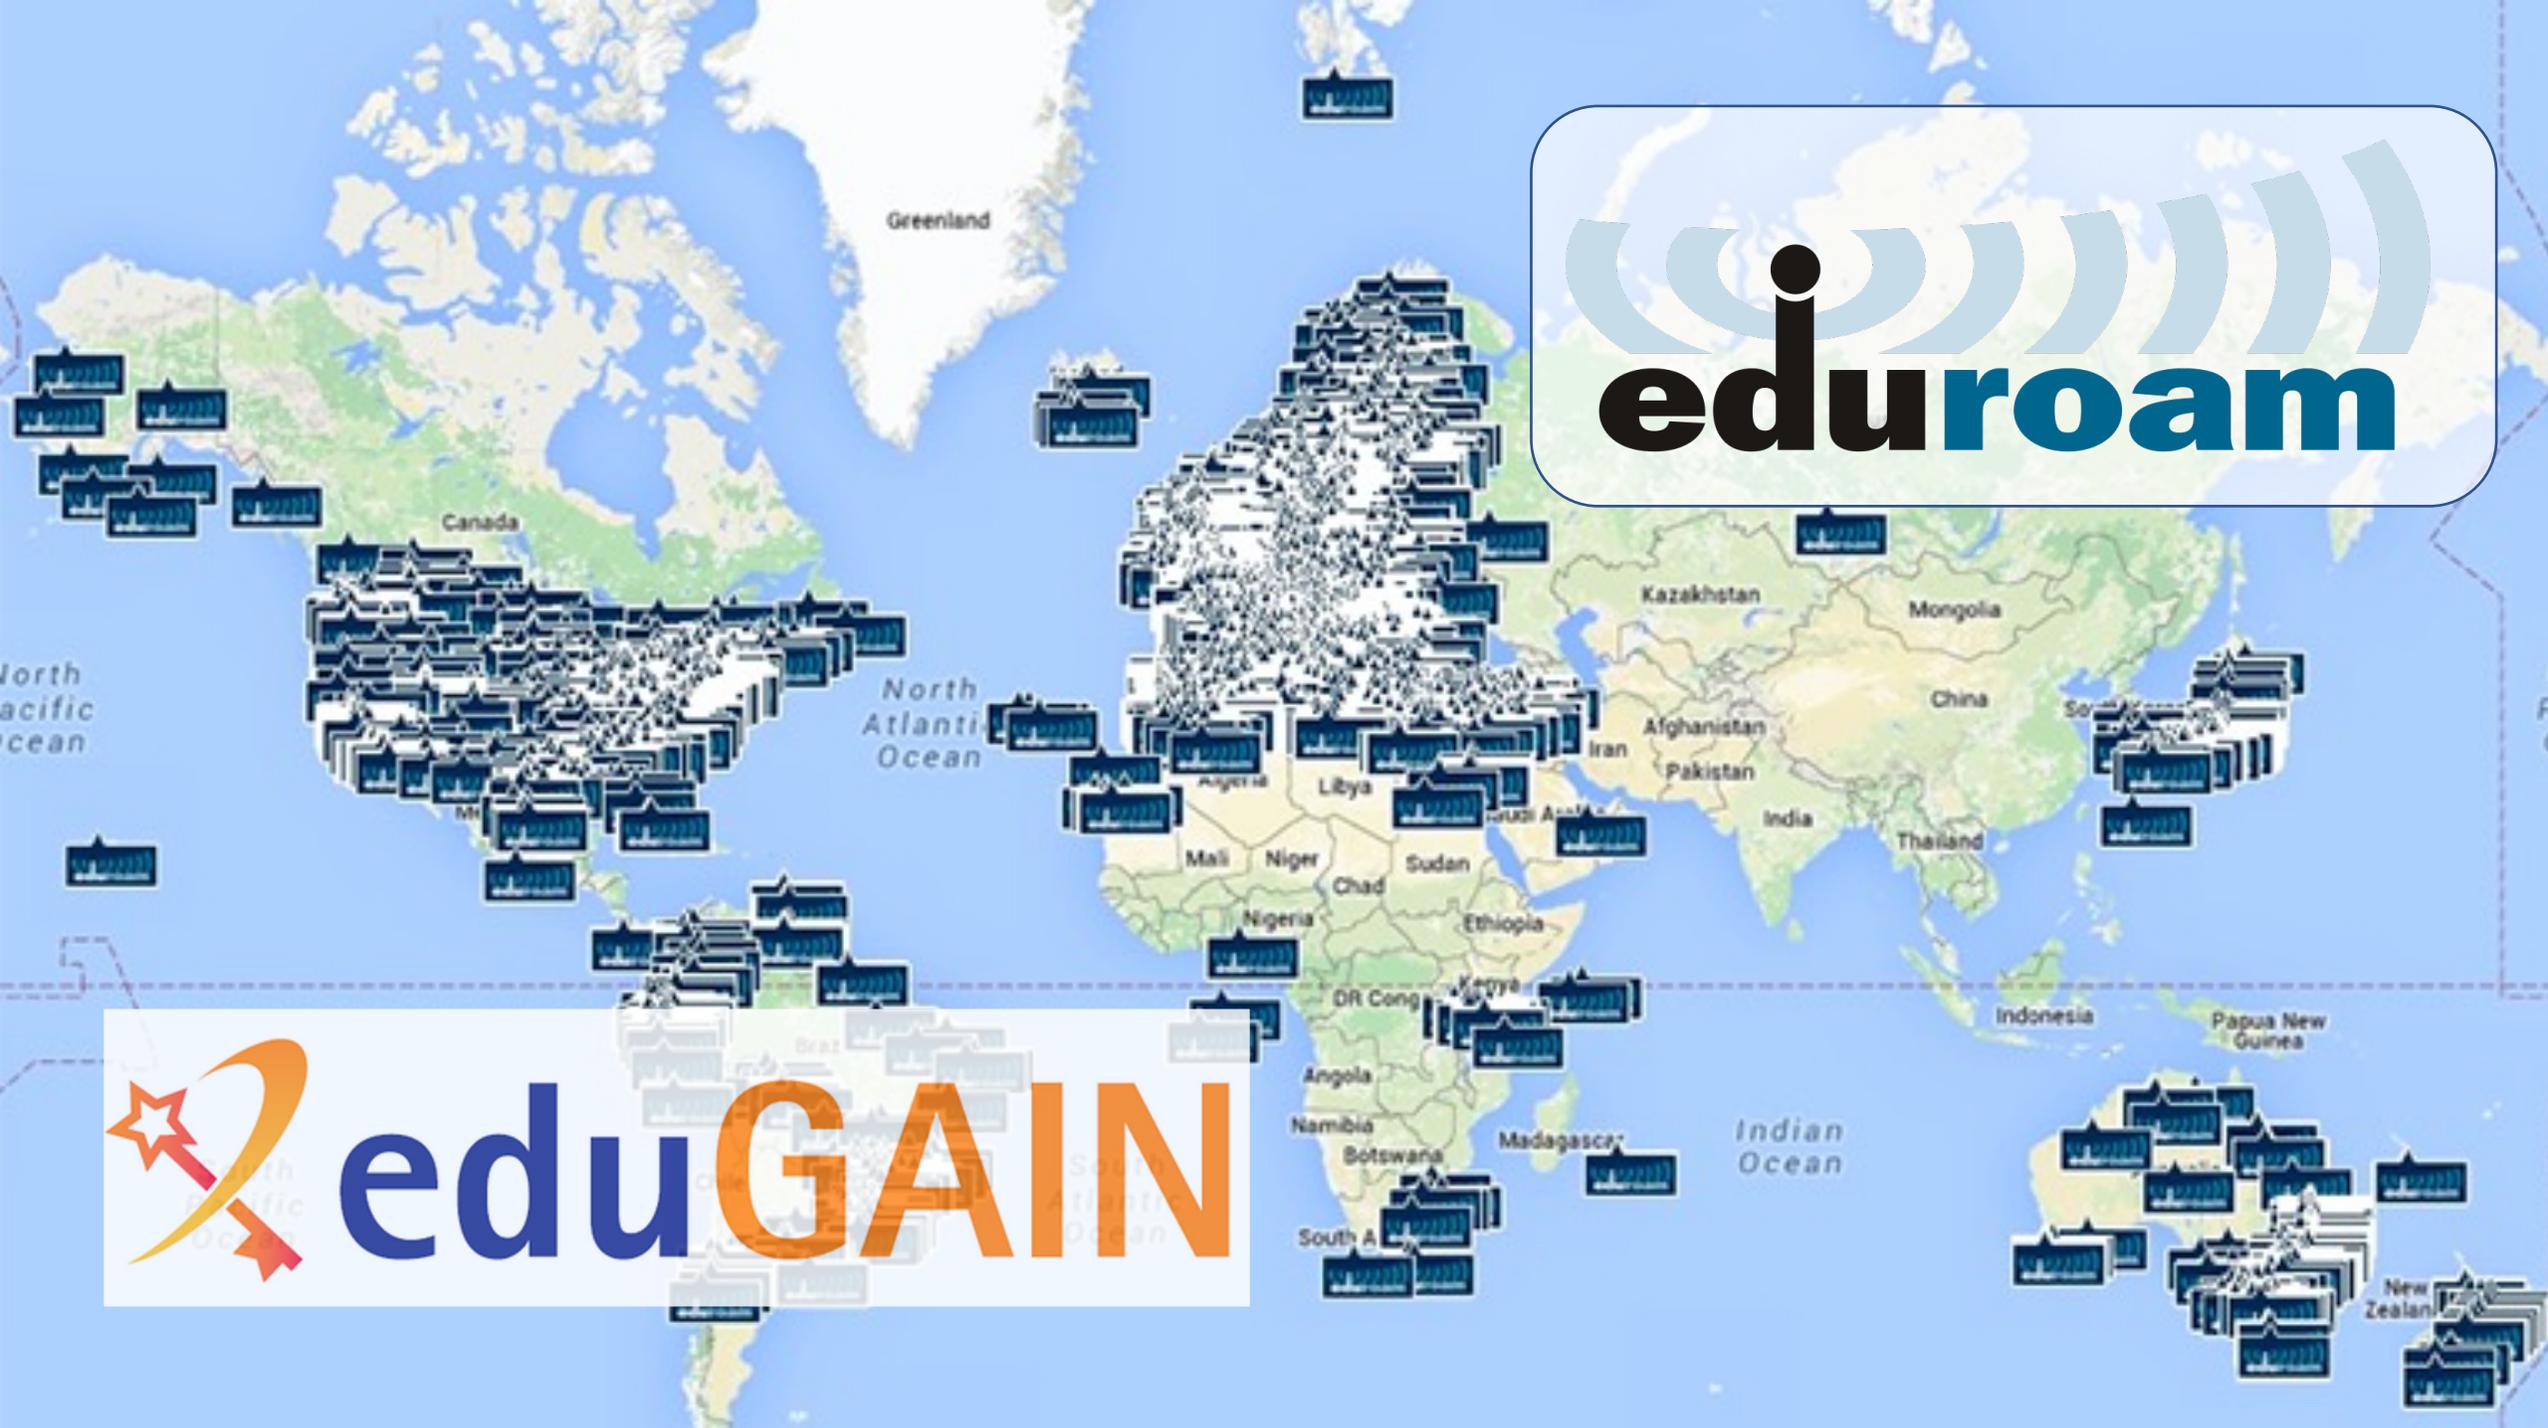


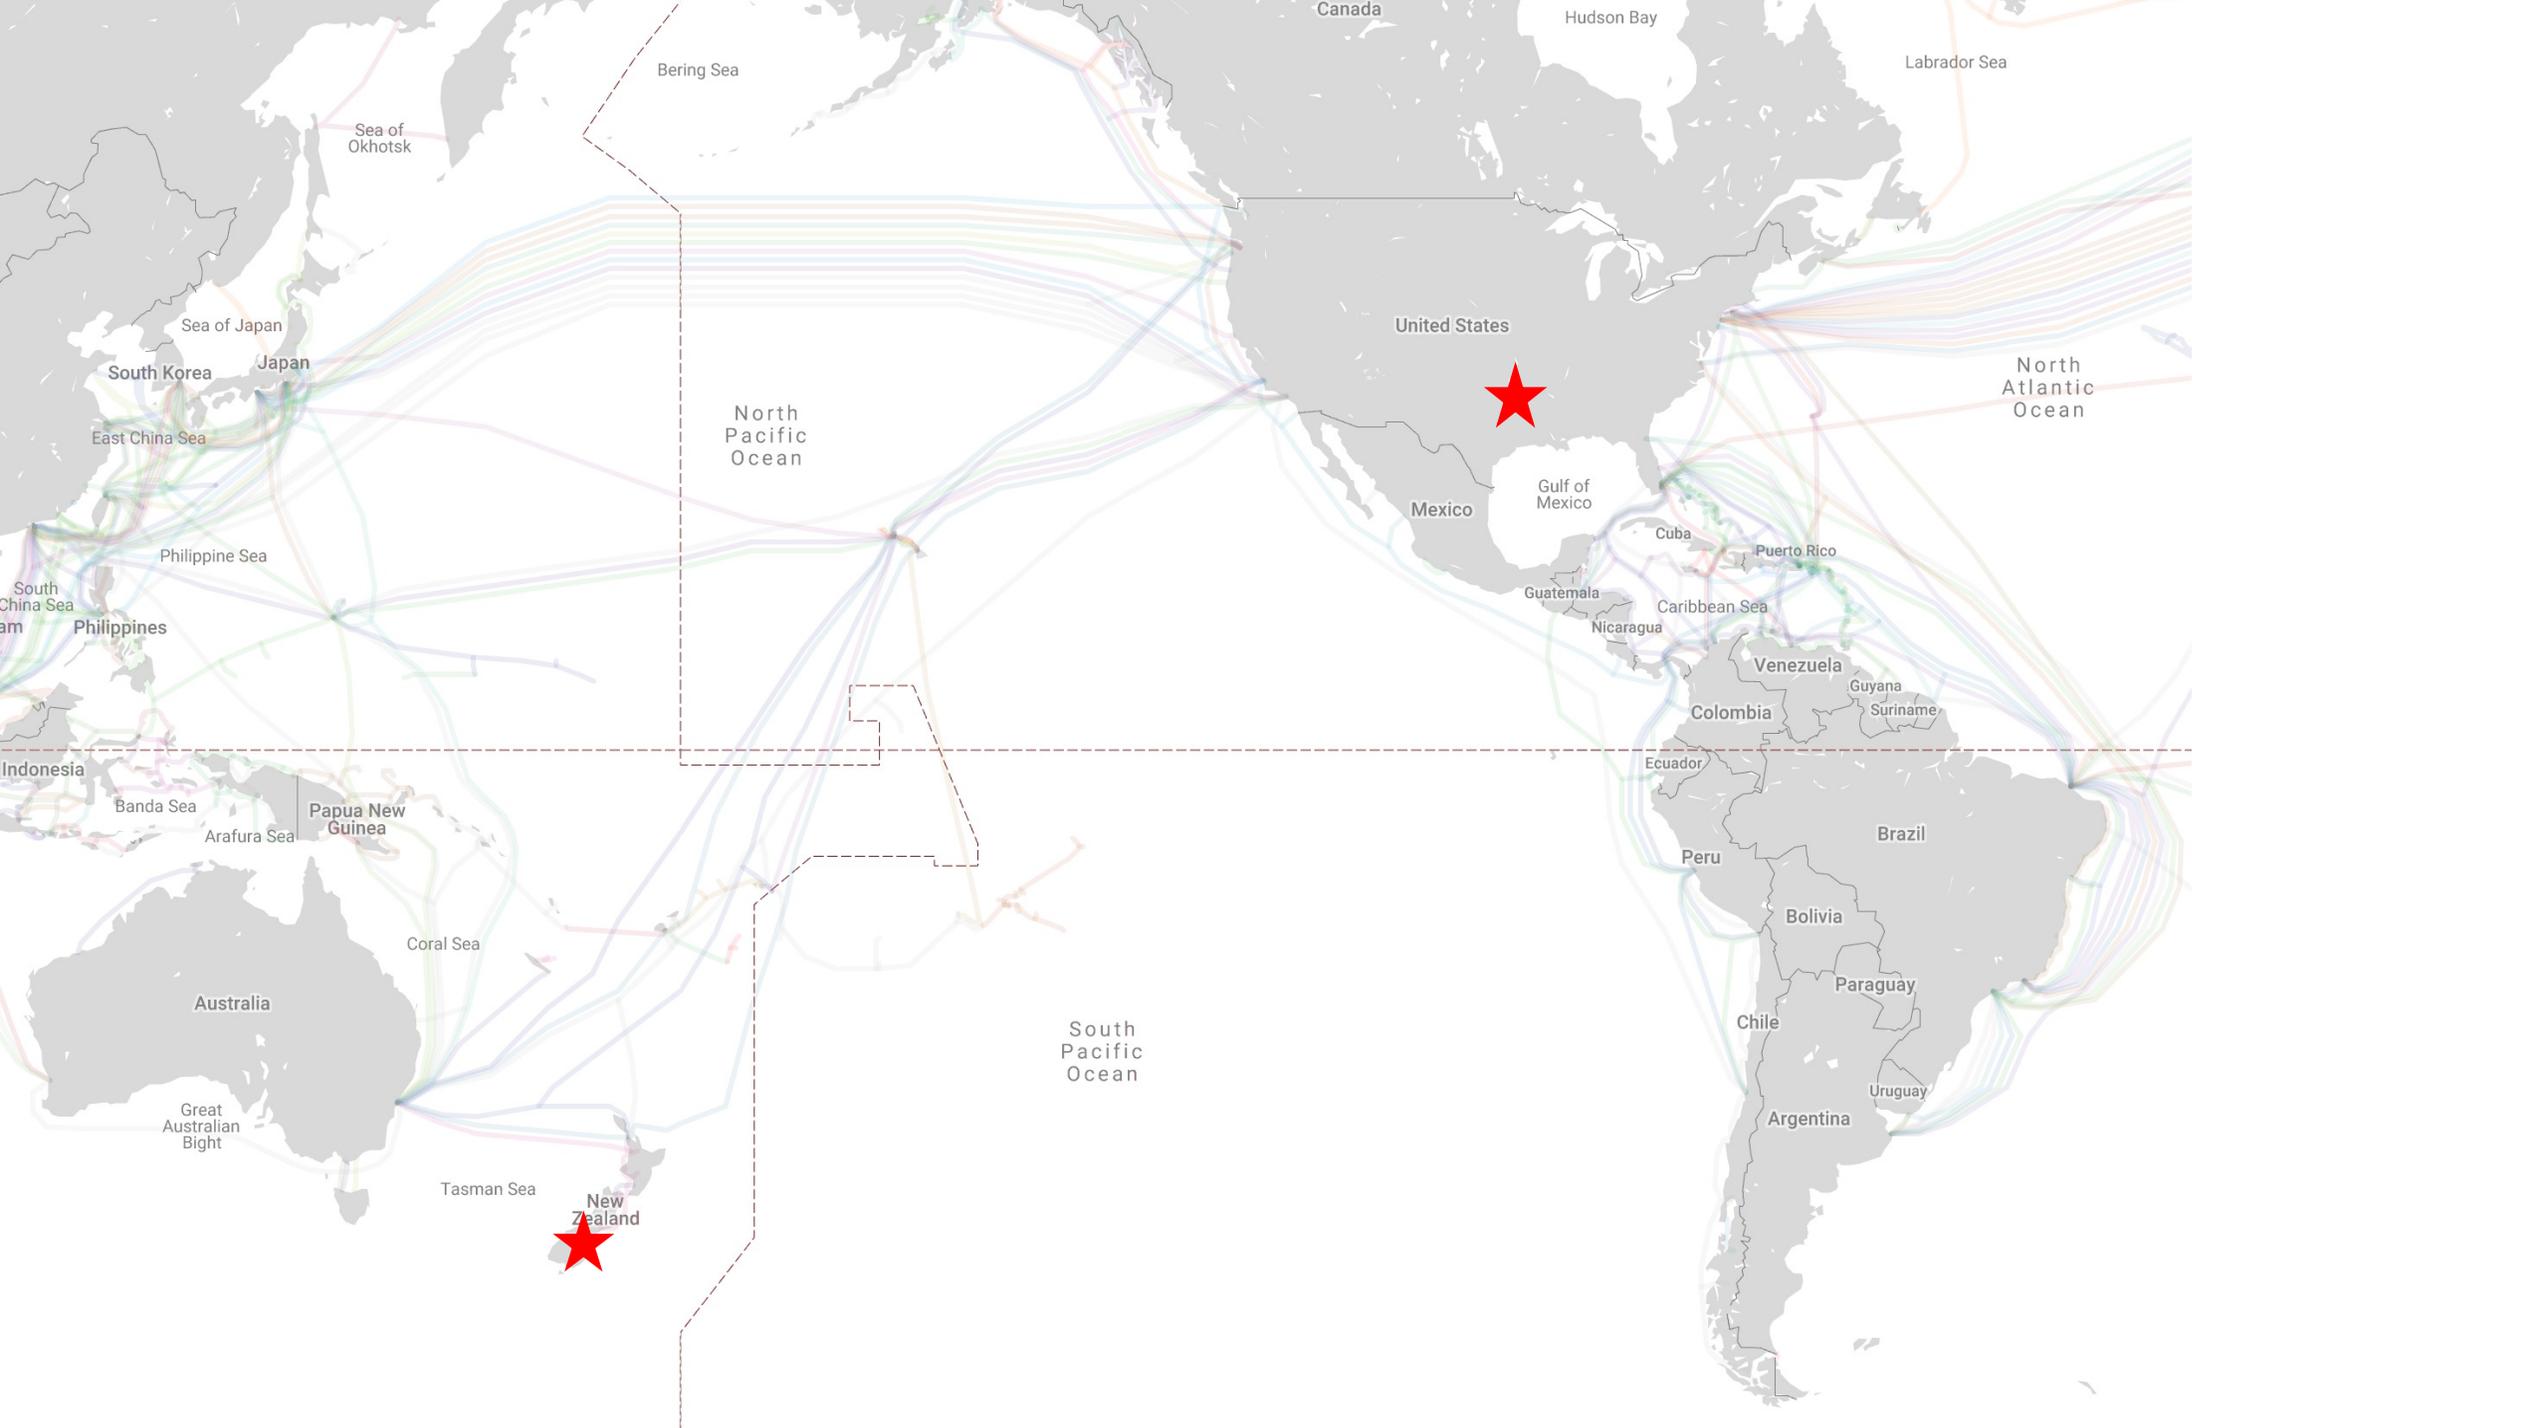


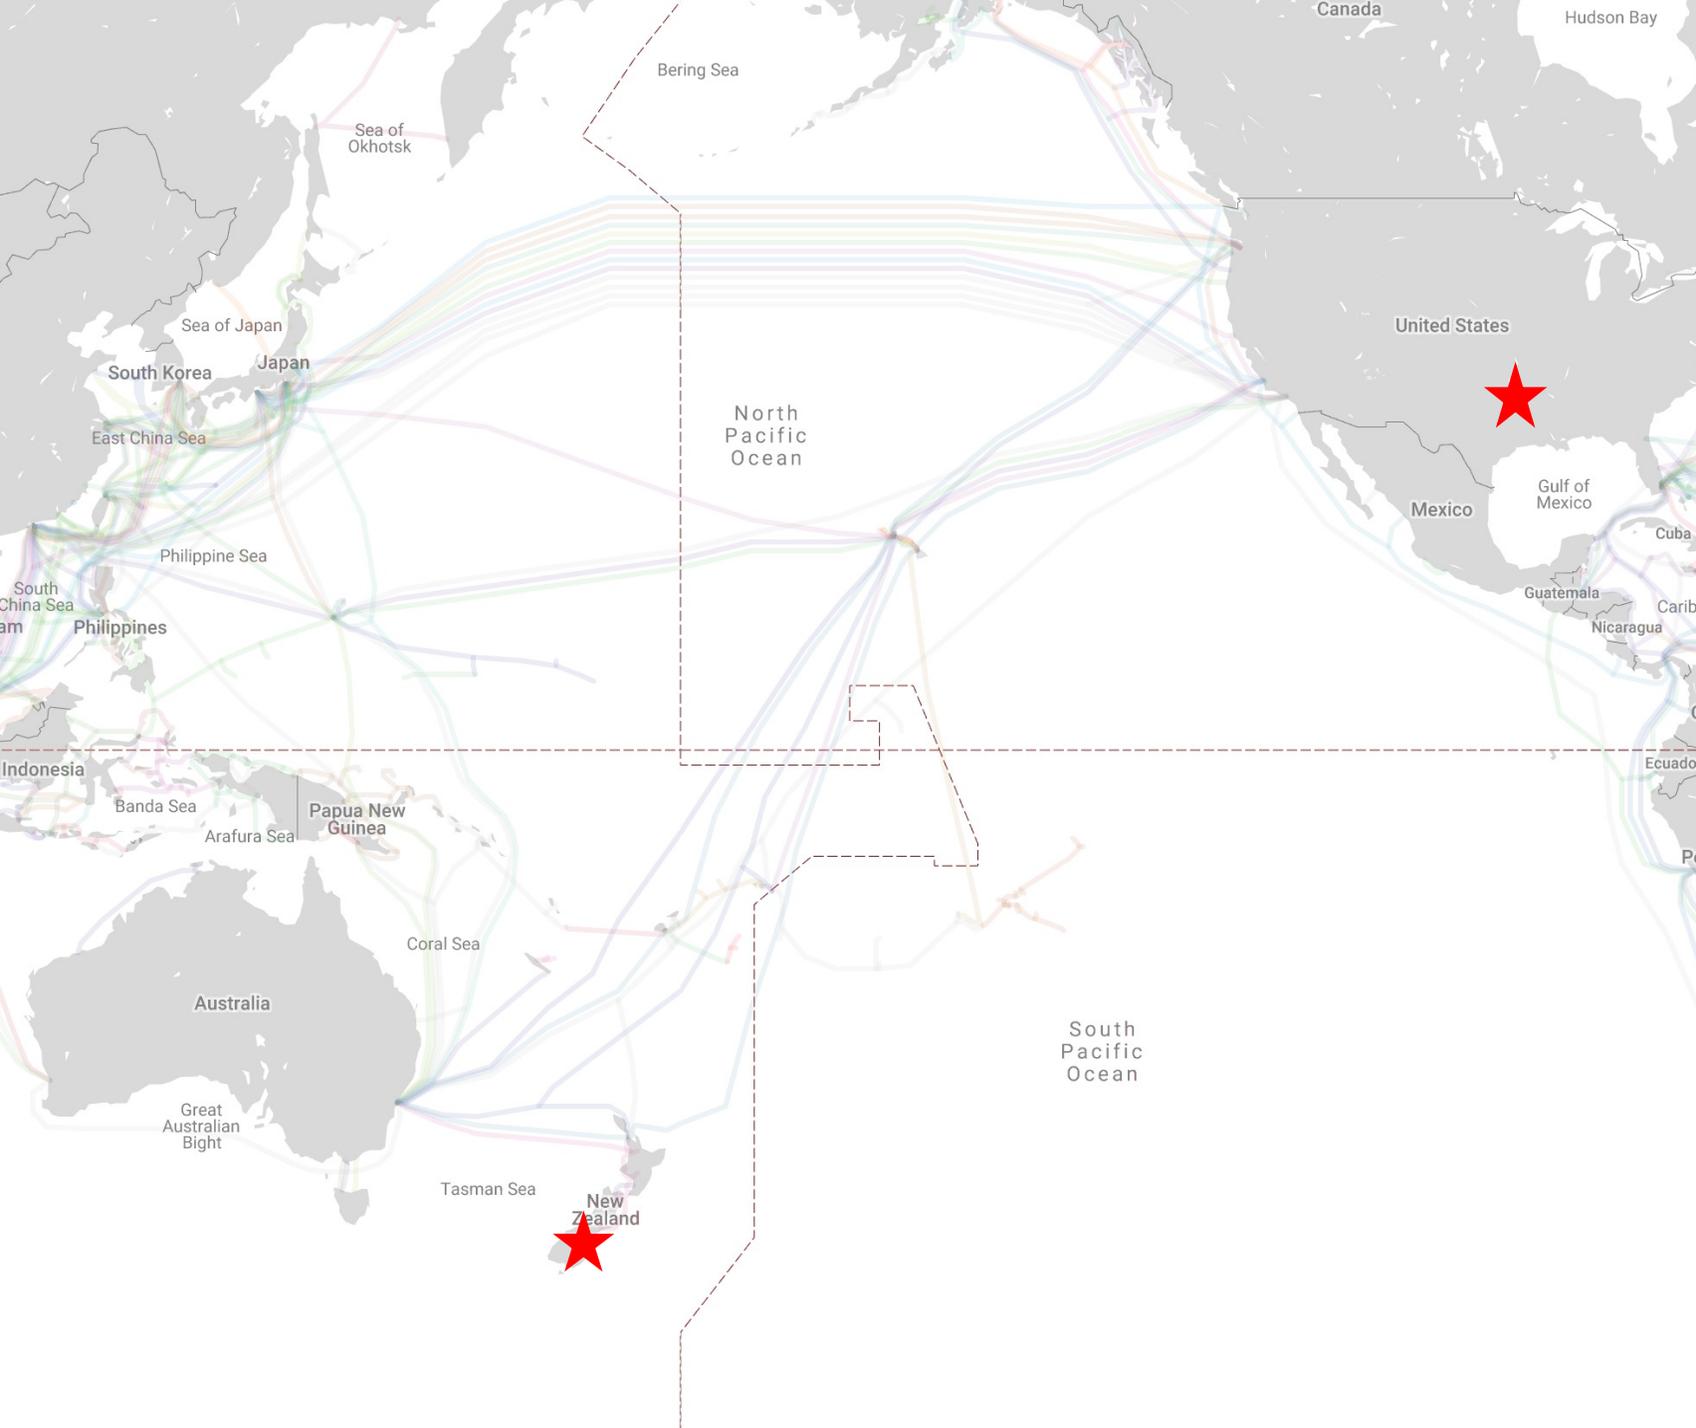
- NorduNet
- Europe/Géant
- UK/Jisc
- Netherlands/SURFnet
- CERN

International community of NRENs enabling collaboration on a global scale!

REANWZ





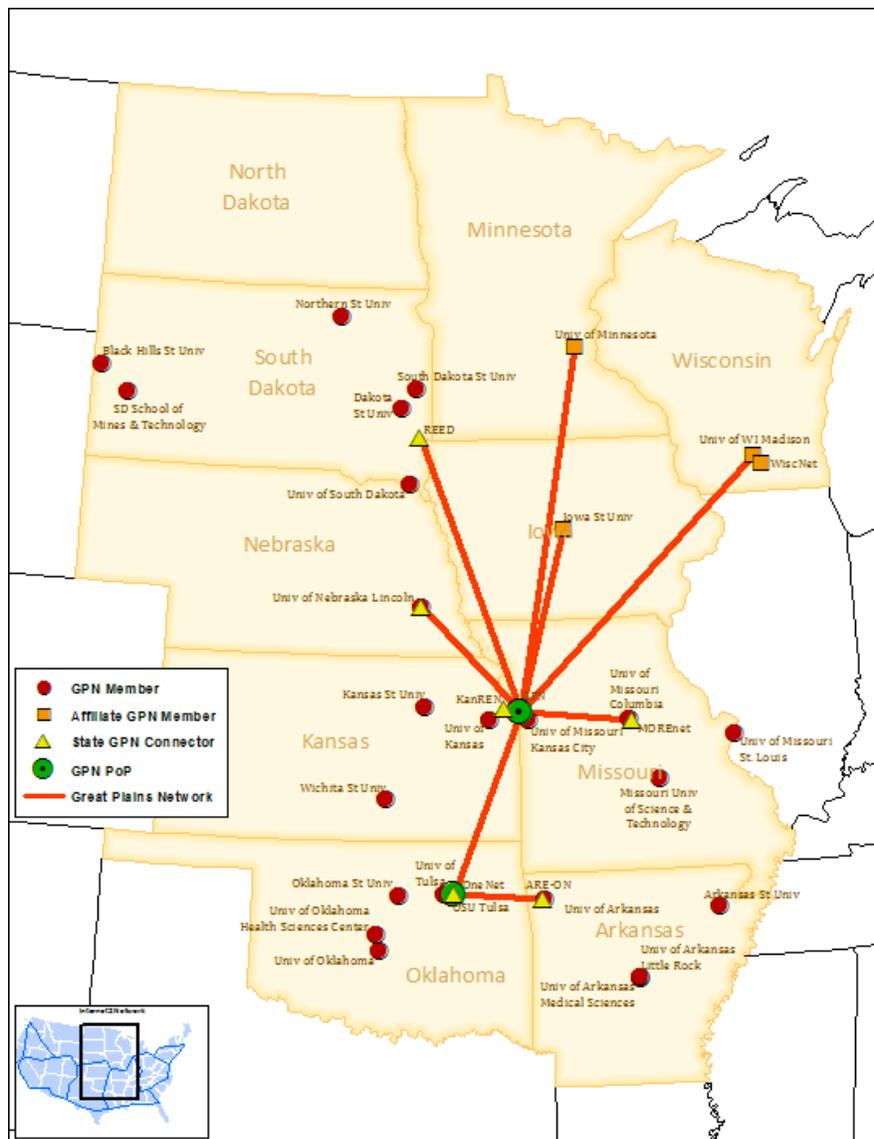




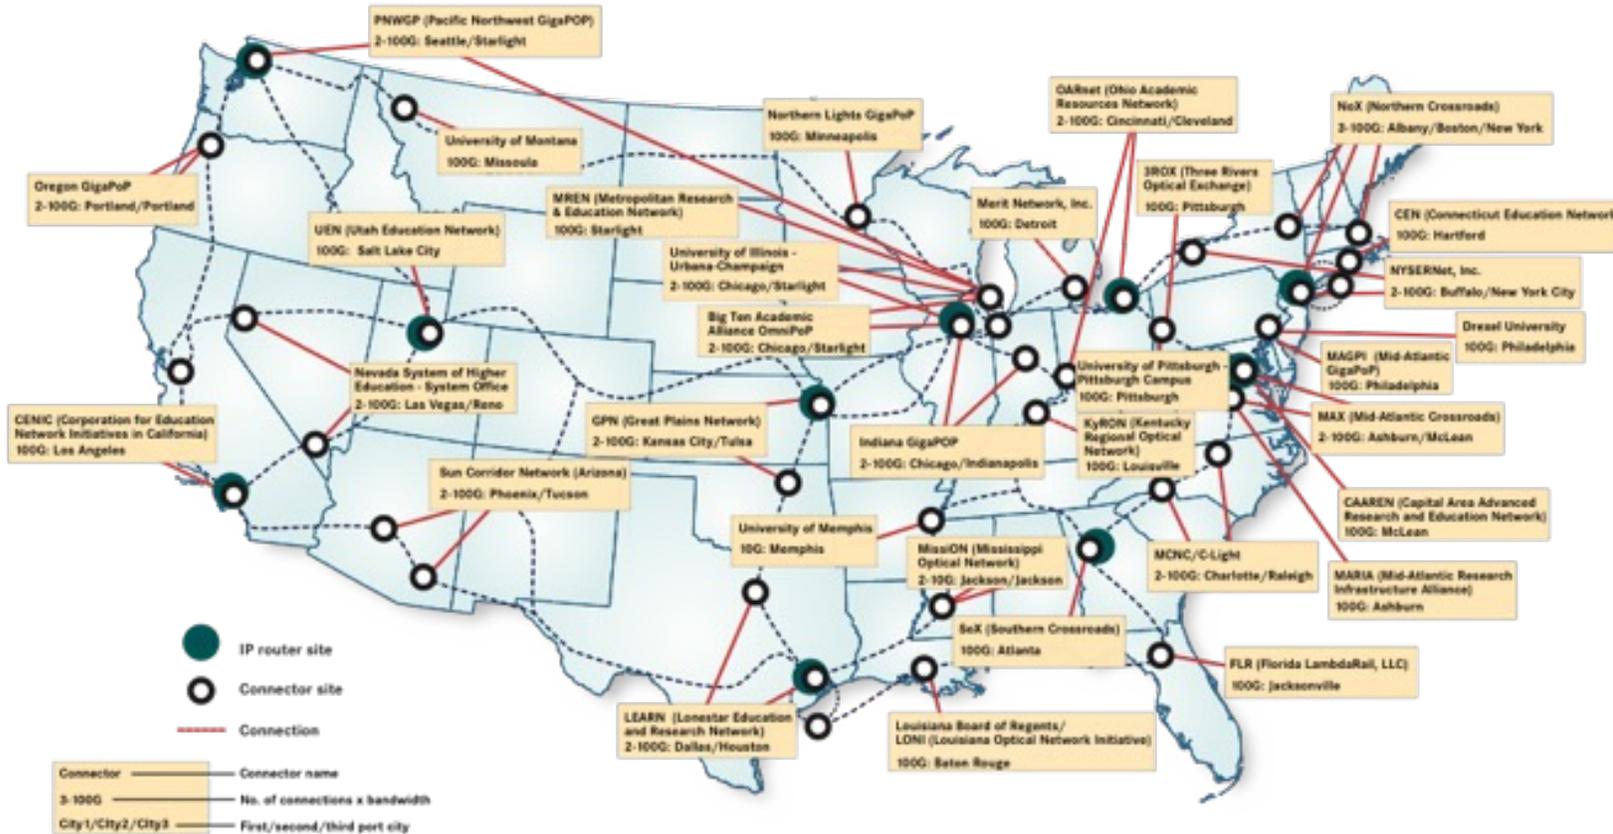
OneNet
Advancing Technology Across Oklahoma

OneNet Points of Presence

- Higher Education Campuses
- Career Technology Centers
- Hospitals
- Public Safety Agencies
- Other



INTERNET2
INTERNET2 NETWORK CONNECTIONS
WWW.INTERNET2.EDU/CONNECTORS - MARCH OF 2017



INTERNET2[®]



NATIONAL & INTERNATIONAL PEERING EXCHANGE

Pacific Wave is a project of CENIC & PNWGP

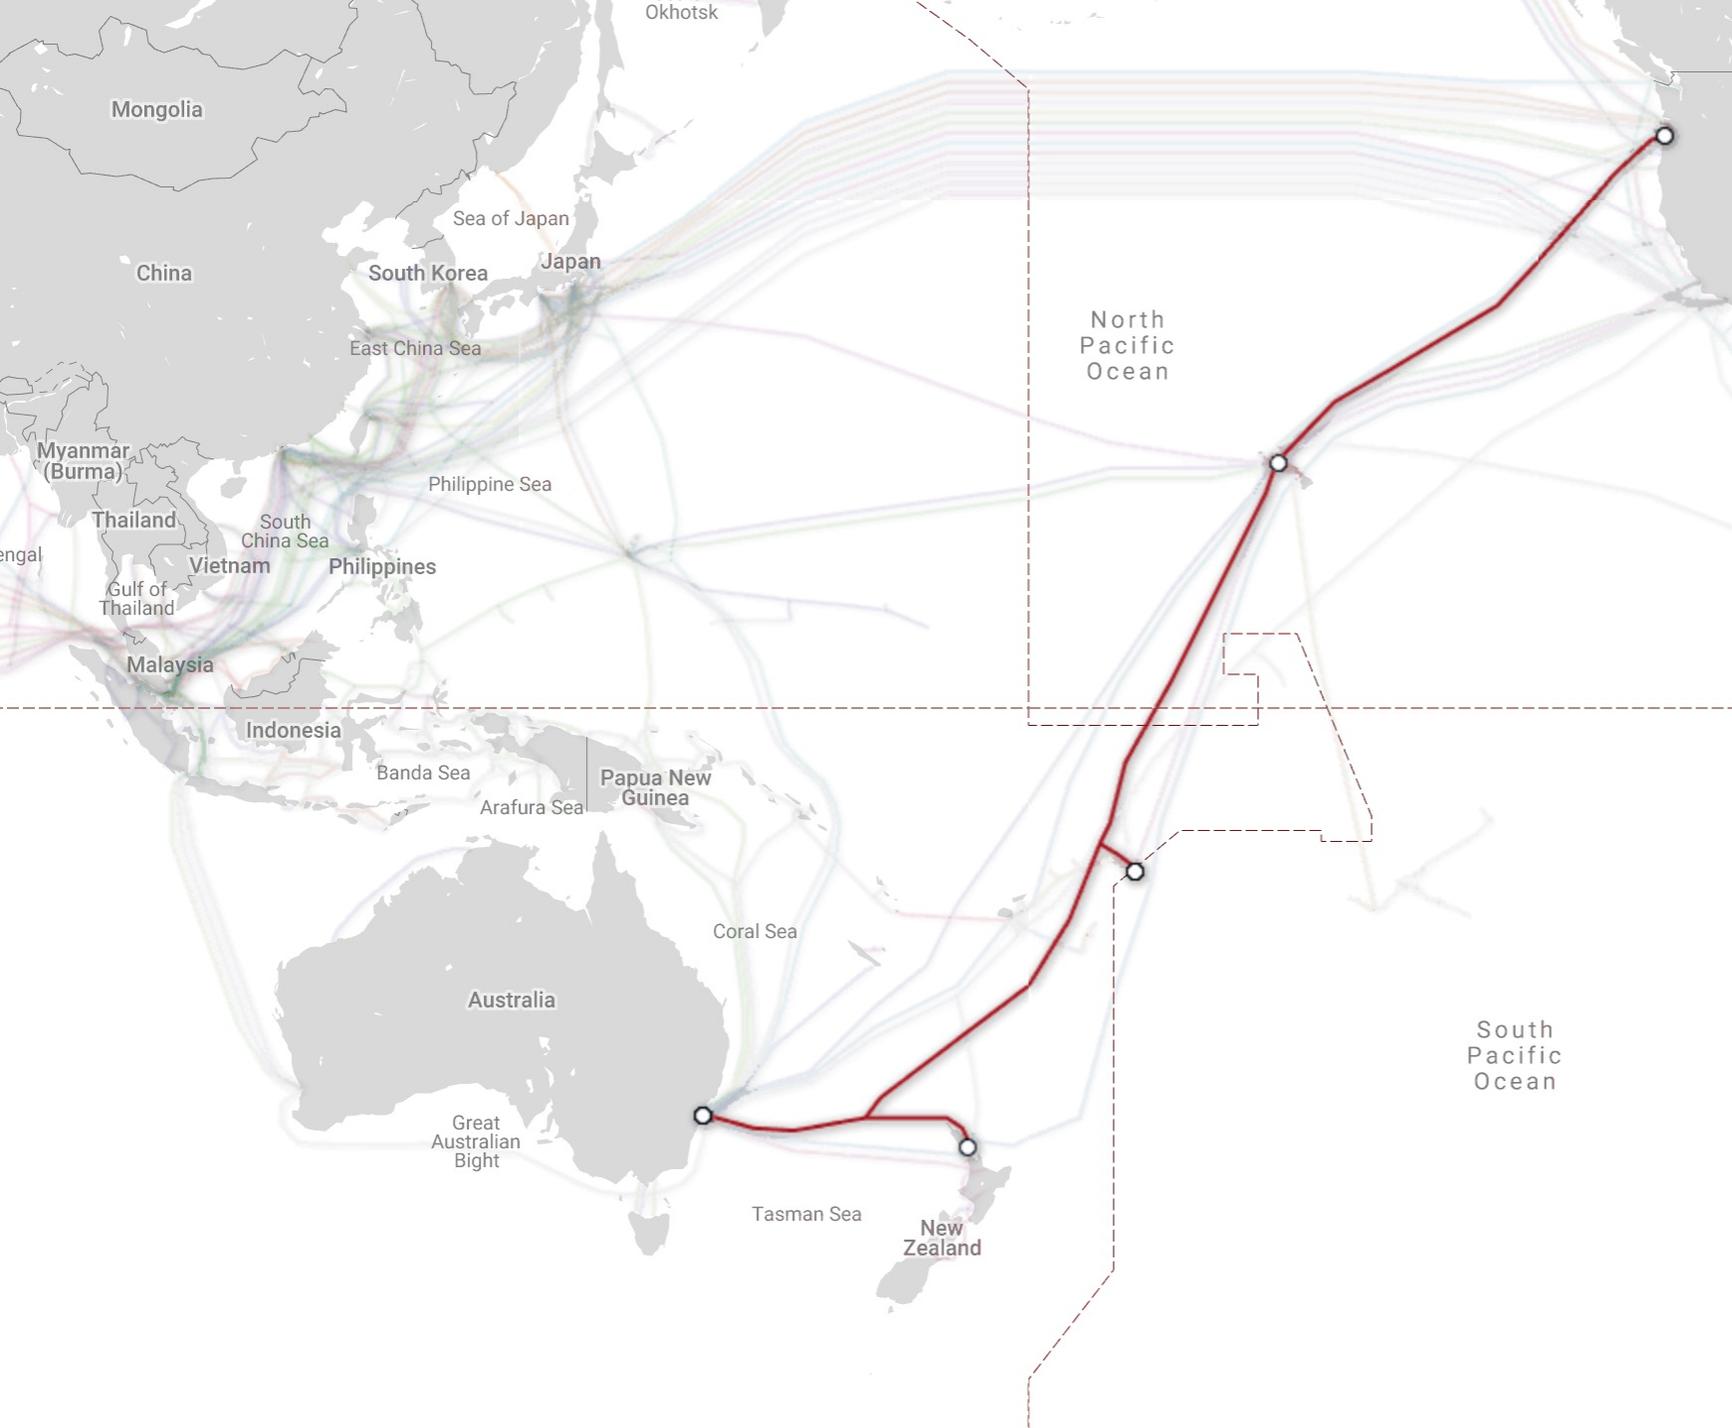


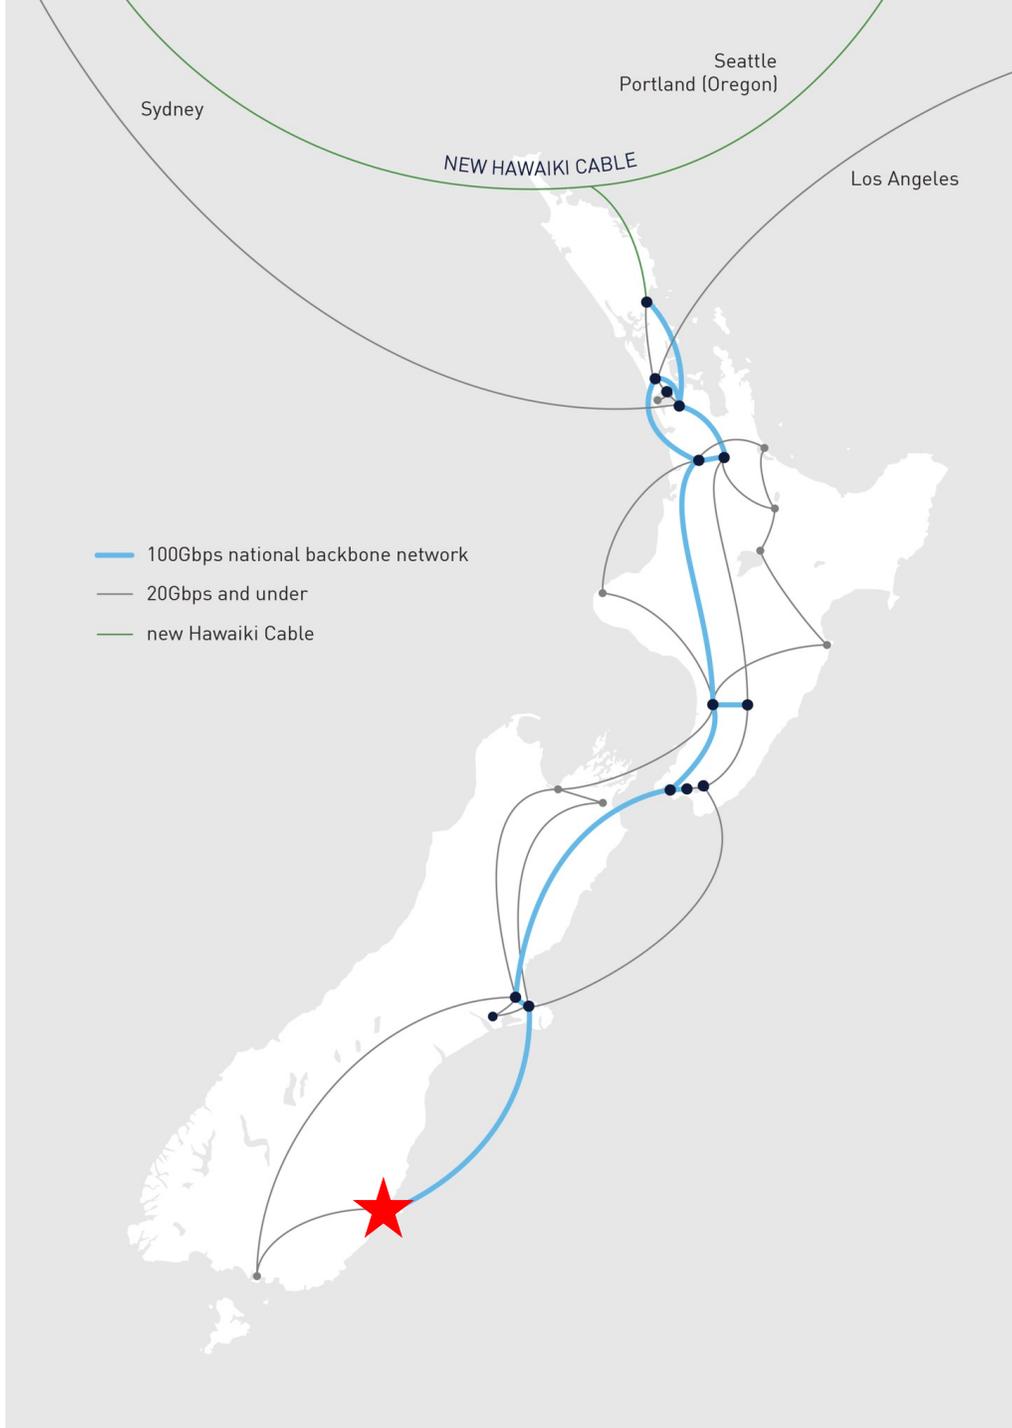
With support from the National Science Foundation



+





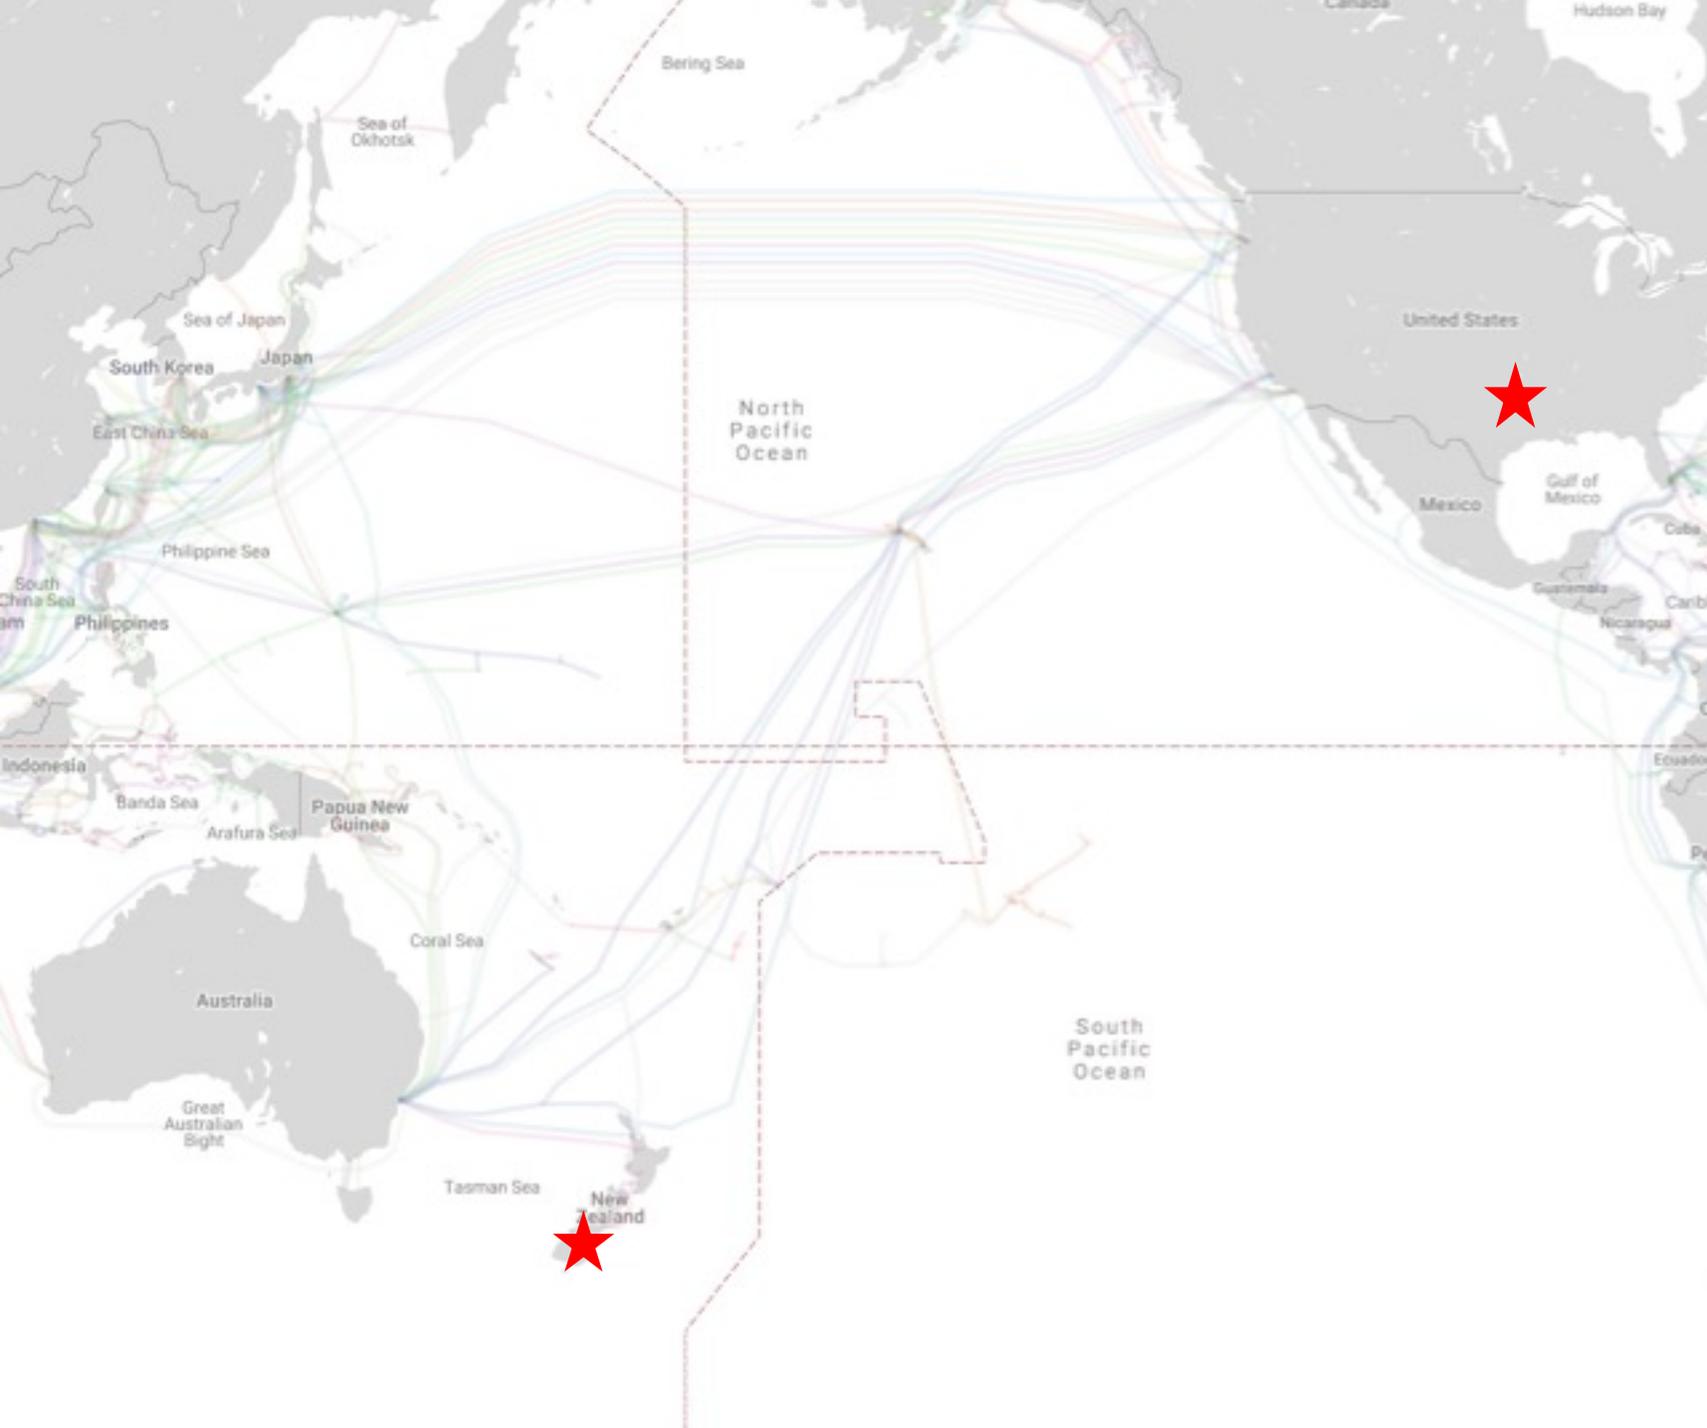


REANNZ



UNIVERSITY
of
OTAGO

Te Whare Wānanga o Otāgo
NEW ZEALAND



So... what does all this mean for me?

- The **requirements** of the data intensive researcher and the **service profile** of the traditional campus computer network (or commercial networks) do not always align!
 - Networks interconnect to networks to networks to networks...
 - Networks get exponentially complex the more connections you have
 - NRENs are at the core of an international community supporting data intensive science!
-
- Research networks are there for YOUR transfers. If its not working, ask for help!

If your data must transit it - its “your” network!

Our goal.

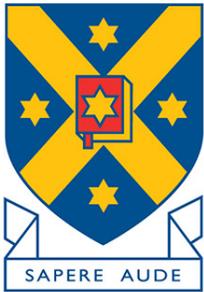


Wallace A. Chase

Head of Department, ITS

wallace.chase@otago.ac.nz

@bmtfr



UNIVERSITY
of
OTAGO

Te Whare Wānanga o Otāgo
NEW ZEALAND

A very special thanks to Matt Younkins at OU, as the raptors are on loan from him!

