

REANNZ

Wednesday 10 February 2016

So you want to send 100GB of data?

T. Charles Yun
eResearchNZ 2016
Queenstown, NZ



So you want to move a BIG data set

- What is “big”
 - Anything that is too big to send as an email attachment
 - Why not just mail a your hard drive?
- The network has changed the way people (scientists, corporate groups, individuals) interact with data
 - The “competition” is already taking advantage of the network
 - Additional funding, reduced costs, improved process, ease-of-use
- This will NOT be a technical talk (xref Ian, no lines of code) (upside: bug free) [and as it turns out, not quite true..., see corrected slide 11]

How “we” think of the network

- Line type (fiber, DSL)
- Line capacity (Gb/s)
- Packet size (jumbo packets, large MTU)
- Congestion (tcp/ip, dropped packets, packet loss)
- Host tuning (kernel, various i/o)
- Application tuning (data staging pipeline, database tuning)
- etc., etc., etc.

The Network Congestion



https://commons.wikimedia.org/wiki/File:Motorcyclists_lane_splitting_in_Bangkok,_Thailand.jpg

Lies, Damn Lies and Statistics...

Fallacy of the station wagon

Never underestimate the bandwidth of a station wagon full of tapes hurtling down the highway.

—Tanenbaum, Andrew S. (1989). Computer Networks. New Jersey: Prentice-Hall. p. 57. ISBN 0-13-166836-6. (taken from Wikipedia)

Lies, Damn Lies and Statistics...

Imagine this scenario...

- Let's say you regularly move data between Auckland and Wellington.
- Distance AKL to WLG: 641 km
- Average drive speed: 80km/h



Lies, Damn Lies and Statistics...

Mazda MX6 Wagon, 2013-2014

- Mazda6 Station Wagon
- Cargo Space: ~500 Liters



<http://www.drive.com.au/it-pro/wagons-v-suv-comparison-test-mazda6-v-mazda-cx5-hyundai-i30-tourer-hyundai-ix35-holden-commodore-sportwagon-v-holden-cotiva7-20140909-10eked>
403-litres


<http://www.carshowroom.com.au/reviews/2012-mazda6-wagon-touring-review-and-road-test/>
519-litres

https://en.wikipedia.org/wiki/File:Japanese_car_accident_blur.jpg

Lies, Damn Lies and Statistics...

LTO-6 Tape

- Linear Tape-Open (2012)
- 2.5TB
- $102.0 \times 105.4 \times 21.5 \text{ mm}$
= 21,501.6 mm
= 0.22l

www.hp.com/go/storagemedia/ultrium

invent

C7972A 400GB
ULTRIUM
LTO 2

https://upload.wikimedia.org/wikipedia/commons/b/be/Lto-4x_hg.jpg

Lies, Damn Lies and Statistics...

Carrying Capacity

- Cargo Space: 500 Liters
- Single Tape Capacity: 2.5TB
- Single Tape Displacement:
 - $102.0 \times 105.4 \times 21.5 \text{ mm} = 21,501.6 \text{ mm} \approx 0.22\text{l}$
- Tapes in Cargo:
 - $500/.22 = 2,272 \approx 2,250$
- Total Data in Cargo:
 - $2,250 * 2.5\text{TB} = 5,625\text{TB}$



ungraciously stolen from: http://www.wallpaperno.com/Humor/funny/minimalistic_funny_swallow_coconut_monty_python_and_the_holy_grail_1600x900_wallpaper_42922/download_1920x1080

Lies, Damn Lies and Statistics...

Fallacy of the station wagon

- 5 Hours to get data in and out of the car:
 - label, sort and box 2,250 tapes
 - load+unload car in AKL and WLG
- 8 Hours to drive AKL-WLG
- $5.6\text{TB}/13\text{ hours} = .43\text{ TB/h}$
 $= 3.44\text{ Tb/h}$
 $= 0.96\text{ Gb/s}$

<http://blog.carchex.com/wp-content/uploads/2014/08/packing-car-6.jpg>

Lies, **Damn Lies** and Statistics...

Fallacy of the station wagon

- 5 Hours to get data in and out of the car:

- label, sort and box 2,250 tapes
- load+unload car in AKL and WLG

- 8 Hours to drive AKL-WLG

- 5.6TB/13 hours = .43 TB/h
= 3.44 Tb/h
= 0.96 Gb/s

5.6TB? derp, that was **5,600TB**. Apologies for getting the math wrong... And belated thanks to the audience for kindly pointing out the mistake

<http://blog.carchex.com/wp-content/uploads/2014/08/packing-car-6.jpg>

Fallacies: corrected, expanded, justified*

- Write data to and from all tapes (or, buying back 3 orders of magnitude error...):
 - write, label, box, read—total 1 hour
 - 2,250 tapes * 1 hours/tape = 2,250 hours
- 5 Hours to get data in and out of the car
- 8 Hours to drive AKL-WLG
- total time: $2250 + 5 + 8 = 2250$ hours

- $5,600\text{TB}/2250 \text{ hours} \sim 2.5 \text{ TB/h} = 20 \text{ Tb/h} = 5.5 \text{ Gb/s}$

* hopefully without errors this time around...

Lies, Damn Lies and Statistics...

Packet Loss

And remember, packet loss in the estate wagon scenario is a pretty big deal

https://en.wikipedia.org/wiki/File:Japanese_car_accident_blur.jpg

Lies, Damn Lies and Statistics...

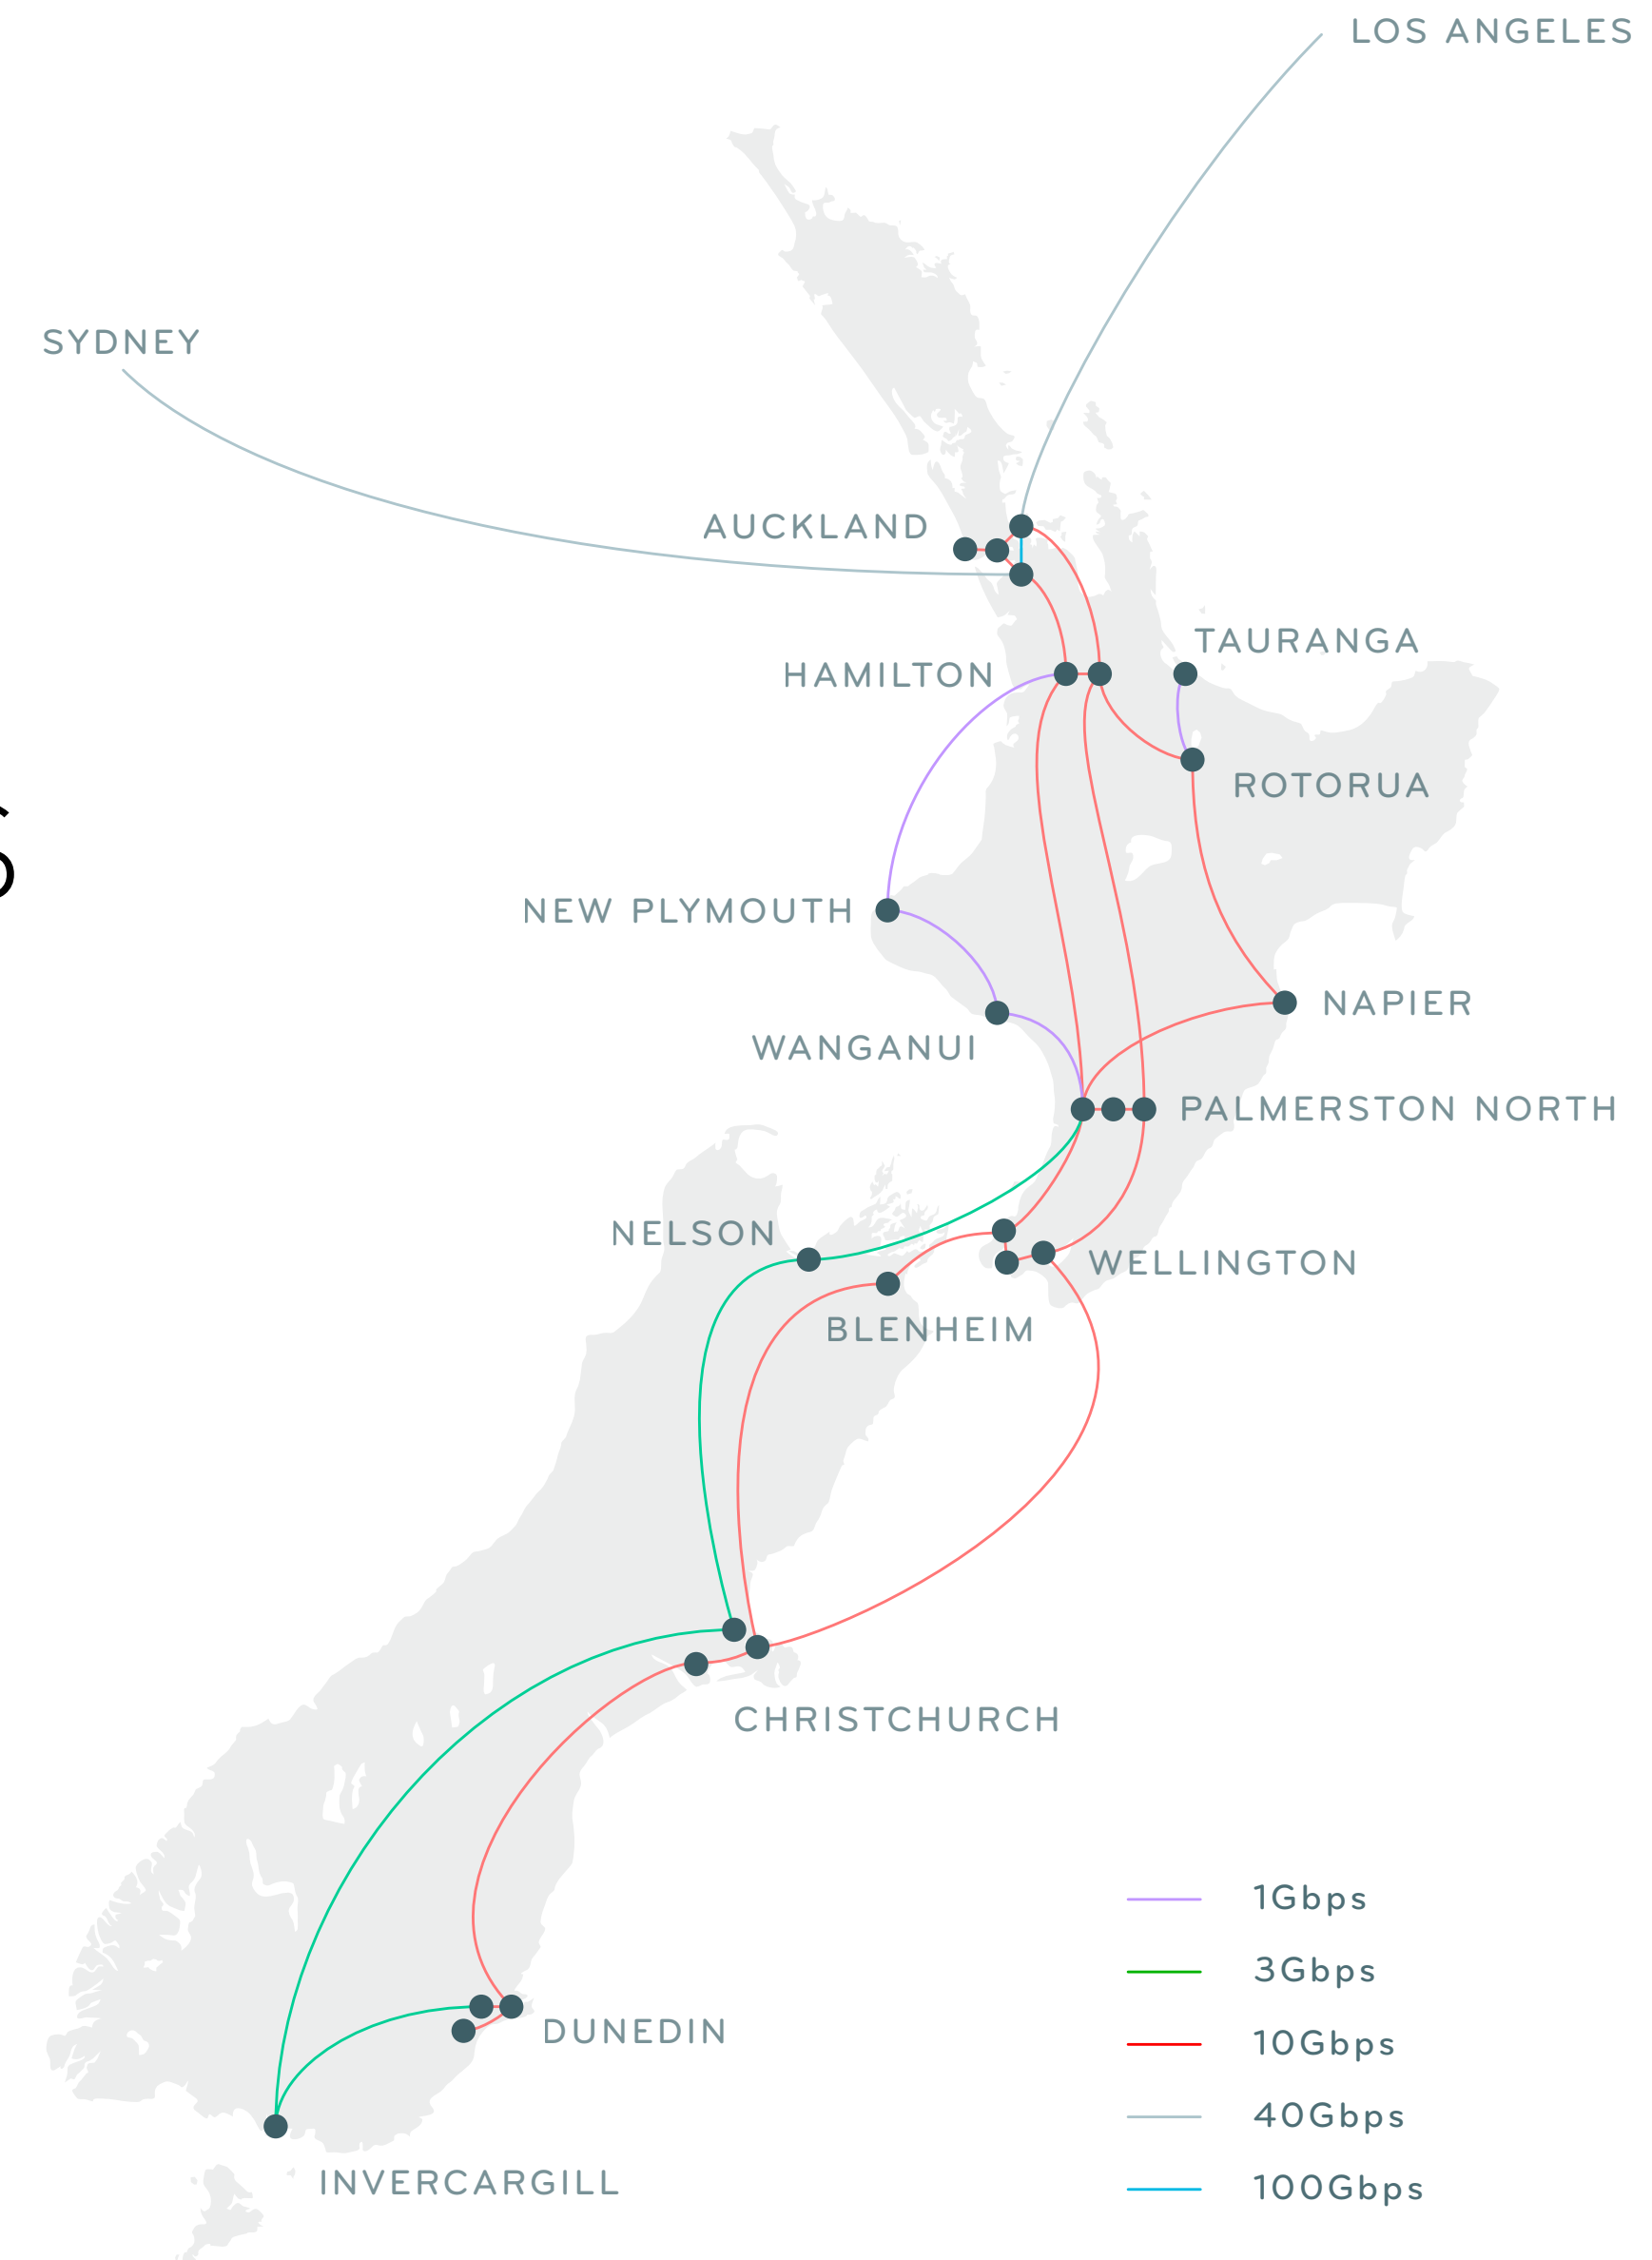
Are you happy with “good enough”

- If you could get 10x improvement in the precision of your scientific equipment by “reading the manual”, would you follow up?
 - If you could stream data continuously, would you even worry about storing files **and then** moving them?
- 1 Gb/s sounds nice
- You should be seeing 10 Gb/s
- We are planning for 100Gb/s
- Everything you need to do better is already in place

Lies, Damn Lies and Statistics...

REANNZ Network

- 10Gb/s ladder design
 - soon: upgrade to 100Gb/s
- International connections to AU and US
 - 40Gb/s R&E
 - ~8Gb/s commodity



Tools

100GB?!

- As a networking organisation, we are interested in **moving** data
 - It is about the network
 - And every problem is a nail
- ...you may want the data to sit still (storage)



<https://upload.wikimedia.org/wikipedia/commons/8/84/Claw-hammer.jpg>

Tools

Globus

- Last year Kyle Chard spoke here about the Globus tools
 - <https://www.globus.org/>
 - <https://www.globus.org/file-transfer>
- Ian Foster, founder of the Globus project, spoke in the previous session on “Building the Modern Research Data Portal”
- Web, CLI and programmatic interface into robust data sharing tools



Tools

Globus even works for New Zealand!

The screenshot shows the Globus website homepage. A red circle highlights a world map in the background, with a red dot indicating New Zealand. A red line points from the text 'Globus even works for New Zealand!' to this dot. The website features a navigation bar with links to Products, News, About, Support, and Log In. The main content area includes sections for 'Fast, Reliable, Secure File Transfer', 'File Sharing Made Simple', and 'UPCOMING EVENTS'. The 'UPCOMING EVENTS' section lists several events, including 'Globus Auth for X science gateway d Webinar' and 'Bio IT World 2016'. A red circle also highlights the 'UPCOMING EVENTS' section.

Research data management x

https://www.globus.org

RWeathermap RData nms rGit rNewsletter rSNOW abs Koordinates The Org Manual Google Analytics Off S2 Sage2 New Tab

globus

Products News About Support Log In

make it easy to move, manage, and share big data.

infrastructure, while providing excellent ease-of-use for your researchers.

greatest inspiration and we love nice things about Globus.

LEARN MORE >

GET STARTED >

LEARN MORE >

GLOBUS PROVIDER PLANS >

USER QUOTES >

CASE STUDIES >

Fast, Reliable, Secure File Transfer

Move files between your laptop, lab server, research computing center, national supercomputing facility, or any other storage system, using just a browser.

LEARN MORE ABOUT FILE TRANSFER WITH GLOBUS >

UPCOMING EVENTS

February 12, 2016

Globus Auth for X science gateway d Webinar

April 5, 2016 to April 7, 2016

Bio IT World 2016

Seaport World Trade Center

April 20, 2016 to April 22, 2016

GlobusWorld 2016

Hotel Allegro, Chicago, IL

VIEW MORE EVENTS >

File Sharing Made Simple

Easily share data with your collaborators, directly from your own storage. All they need is a free Globus account, and you control what files they can access.

LEARN MORE ABOUT FILE SHARING WITH GLOBUS >

CloudStor

- Today, you will hear David Jericho talk about CloudStor (in about 15 minutes)



- CloudStor: Getting Out of the Way of Researchers
- <https://www.aarnet.edu.au/network-and-services/cloud-services-applications/cloudstor>
- Integrates with Tuakiri
- 10GB of free storage (100G if you are AU), unlimited File Sender (web based upload)
- Free, web-based tool

ScienceDMZ

- A “design pattern” from ESnet (xref Inder Monga’s talk yesterday)
- A tool to help separate research and corporate data by explicitly working with the participating parties (source, destination && {LMW}AN networks
- Several in place here in New Zealand today

Tools

Aspera

- Sean Curran spoke earlier today
- Used in Australia through a national site license
- Being tested by radio astronomers who are interested in spending money on **their** science, not file transfer protocols
- And yes, this is a big enough problem that you can even create a business around a solution



Getting back to the 100GB of data...

- We've spoken a bit about wanting to move 100GB
 - We are against mailing hard drives
 - We have shown you some tools as alternatives
- But how did we get here?
 - Why do you have 100GB all at once

E2E Testing

- What level of performance do you ***think*** you should be seeing?
- End-to-End Testing
 - iperf, perfSONAR
 - ...tested your disk array?
- did you happen to launch your transfer via a poorly written script?

Change Expectations

How long will it take?

Data set size				
10PB	1,333.33 Tbps	266.67 Tbps	66.67 Tbps	22.22 Tbps
1PB	133.33 Tbps	26.67 Tbps	6.67 Tbps	2.22 Tbps
100TB	13.33 Tbps	2.67 Tbps	666.67 Gbps	222.22 Gbps
10TB	1.33 Tbps	266.67 Gbps	66.67 Gbps	22.22 Gbps
1TB	133.33 Gbps	26.67 Gbps	6.67 Gbps	2.22 Gbps
100GB	13.33 Gbps	2.67 Gbps	666.67 Mbps	222.22 Mbps
10GB	1.33 Gbps	266.67 Mbps	66.67 Mbps	22.22 Mbps
1GB	133.33 Mbps	26.67 Mbps	6.67 Mbps	2.22 Mbps
100MB	13.33 Mbps	2.67 Mbps	0.67 Mbps	0.22 Mbps
	1 Minute	5 Minutes	20 Minutes	1 Hour
Time to transfer				

<http://fasterdata.es.net/home/requirements-and-expectations>

Change Expectations

Planning

- Pipeline
 - Can you change your pipeline to stream the data constantly?
 - Instead of bulk delivery of 100GB, constantly stream data
 - Spread the impact on the network over a longer period
- Where do you really need the data?
 - Why is it moving in the first place? Can it be analysed while it is moving so that you do not need to move it again?

Raising Expectations

- Not only should New Zealand researchers have access to the tools and infrastructure that their international colleagues have, but they should have the expectation that the tools and infrastructures will work well.

Implementing Tools

- REANNZ is running the backbone and increasingly assisting with building the on-ramps (xref Rob Elshire's plenary talk from earlier today)
- REANNZ must rely on the users to build the (fast) cars that travel on the road
 - eResearchNZ is one way to execute outreach
 - We work closely with IT Departments to improve on-campus network performance

Change Expectations

Tuned System



https://upload.wikimedia.org/wikipedia/commons/6/61/Silverstone_2011_starting_grid.jpg

Questions?

T. Charles Yun

tcyun@reannz.co.nz

<http://www.reannz.co.nz/>
@REANNZnews





—fin—