

# Learnings from running a free public cloud for 10 Years

Micheal Jones



# Agenda

- Context
  - What is the Rapid Access Cloud?
  - Who is Cybera?
- What did we learn?
- Q&A

# Slide apology

- My slides aren't pretty
- Slide link if you're curious
  - Not a rick roll :-)





# What is the Rapid Access Cloud?



# What is Rapid Access Cloud?

- Started in 2012 - "What happens if we give folks free compute"?
- OpenStack based cloud free for Albertans to use
- For **non-production** workloads and education uses



# RAPIDACCESS

C L O U D

Log in

**User Name**

**Password**

**Region**

[About RAC](#) | [Updates](#) | [Help](#)

Connect

Project ▾

Compute ▾

Project / Compute / Overview

Overview

# Overview

Instances

Volumes

Images

Key Pairs

API Access

Network &gt;

Orchestration &gt;

DNS &gt;

Object Store &gt;

RAC &gt;

Identity &gt;

## Limit Summary



Instances

Used 6 of 100



VCPUs

Used 11 of 150



RAM

Used 15GB of 1000GB



Floating IPs

Used 3 of 10



Security Groups

Used 13 of 20



Volumes

Used 10 of 75



Volume Storage

Used 635GB of 1.5TB

## Usage Summary

# History in a Hurry

- 2012 - Started with a controller from Nebula and a handful of compute nodes
- 2014 - Recreated with our own deployed controllers and nodes in both Edmonton and Calgary
- 2016 - In conjunction with CANARIE, started incorporating GPUs
- Has grown to ~1700 cores, 12TB of RAM, and ~90 TB of object storage space available to use.





But who is Cybera?



# HOW CYBERA SUPPORTS ALBERTA'S PUBLIC SECTOR

COLLABORATING ON IT TOOLS AND SERVICES



ADVOCATING FOR ACCESS



ENABLING NEXT-GENERATION SKILLS



DATA VAULT



SECURING NETWORK DATA



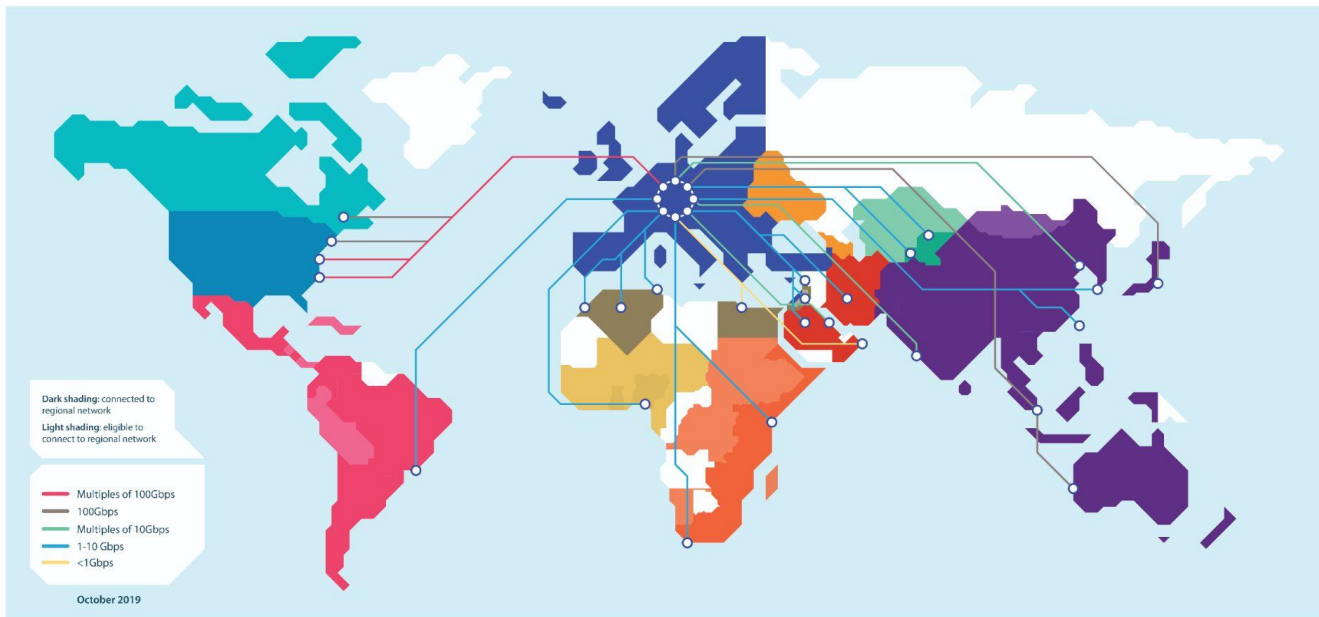
CONNECTING HIGH-SPEED NETWORKS



# What We Do



# AT THE HEART OF GLOBAL RESEARCH AND EDUCATION NETWORKING



Canada & USA	Latin America	Europe	North Africa & Eastern Mediterranean	West & Central Africa	Eastern & Southern Africa	Central Asia	Asia-Pacific	Other R&E Networks



This project is part of the GÉANT OpenRt-Cross Agreement (GSA 3.0), 06/2016, that has received funding from the European Union's Horizon 2020 research and innovation programme under the GÉANT 2020 framework Partnership Agreement (no. 94948). The content of this document is the sole responsibility of GÉANT and can under no circumstances be regarded as reflecting the position of the European Union.

[geant.org](http://geant.org)

# Also Other Services

- Shared procurement
- Internet Buying Group for K-12 school boards in the province
- Data Science Training and Programs
- Callysto (Jupyter Notebooks and Curriculum for Grade 5-12)



What'd we learn?



# Are you an OpenStack admin or ???

- Dependency Management - Day 2 doesn't mess around.
  - MySQL clusters
  - RabbitMQ clusters
  - Human processes
  - Abstraction management

# Are you an OpenStack admin or ???

- Complex systems tend to... get more complex.
- What would you say to how a failing hard drive in Edmonton led to a Calgary network outage?
  - Root Cause Analysis is useful - but not the end of the story. It's also a useful way to lie to yourself.



# Are you an OpenStack admin or ???

- OpenStack is Python "glue" on top of known technologies
  - Libvirt/KVM
  - iSCSI
  - Linux Bridging
- Metrics and GUI apps are awesome
- strace, tcpdump, sar, and nc get the job done

# IP Addressing

- RAC offers public IPv4 AND IPv6 addresses
- You can never have enough IPv4 addresses and there are not enough.
  - One class of 150 CompSci students can burn through your address pool real fast

# IP Addressing

- IPv6 scares folks - because change is scary.
- IPv6 adoption in Alberta is... not great.
  - TELUS - supposedly on every home connection
  - Shaw/Rogers - some modems but not all
  - Everyone else - mixed bag if it's turned on
  - Universities - very slow
  - Supernet - still doesn't support it at all.

# "Free" is hard

- How do you keep support manageable?
  - Documentation!
  - Expectations
- How do you keep resources available?
- How do you handle that you've removed the market force to encourage less waste?
  - Do you actually see more experimentation with the constraint removed?

# "Free" is hard

- What is the operational tradeoff?
  - Hardware ages. Hardware breaks. Software needs updates.
  - GPUs just keep getting more expensive each year

# Community is priceless

- OpenStack Foundation did a good job fostering and creating communities
  - Mailing Lists
  - IRC
  - Conferences
- Global - we've worked with folks all over the world with it
- One of my favourite sayings from Kubernetes folks was "how do we not repeat OpenStack's mistakes?"

# Containerization is awesome

- Not exactly new in 2014, and definitely not new now BUT
- It's a fantastic approach operationally.
  - Solved a lot of issues of running several Python services along side and avoiding conflicts between versions.
- Embracing CI/CD is "easier" with containers

# Abuse

- Far less common than you may presume
- Far less common != not at all
- Attacks targeting bad configs on the cloud are a much bigger deal.



# What'd we learn?

- Trust in humanity – people will surprise you.
- Trust but verify.
- Design for people – making things secure, private, the default and easy option pays off.

# Importance of Self Hosting - Services

- Commodity level services is a hard sell
  - Email - most of our email is already on Google or Microsoft's servers.
  - Anti-spam and dealing with blacklists is hard.
- Services that are latency sensitive, or key to the business
  - File Sharing (eg. documents, video)

# Importance of Self Hosting - Infrastructure

- Cloud services can get expensive quickly.
  - The tradeoff is real - not needing to manage your hardware and have the elasticity is fantastic. Unless you don't need those features.

# What did we get wrong about self hosting?

- The desire to avoid managing hardware is bigger than we thought
  - People not interested in hardware.
  - Changing dead hard drives is tedious and not fun.
- Underestimating the value of outsourcing "the hard bits"
  - Physical security, data protection certifications

# What things surprised you?

- Minecraft Servers
- OwnCloud

# Network Virtualization

- Just like how physical compute hardware was moved, people want to move network devices as well
- NFV is even more selective about it's location
  - Similar to how compute is best placed close to your data
  - So for us - virtual firewalls for members makes sense.
  
- The network path is everything (YYCIX shoutout)

**"Cloudy" is hard**

<https://xkcd.com/1053/>

# "Cloudy" is Hard

- Our users tend to be very new to the "cloud"
- Word of mouth or post secondary students/researchers
- Hardware Requirements != Cloud Requirements
  
- How do we better help those new to cloud, or even new to using a hosted services use them more effectively?



# Fads

- We'd love to better identify passing fads versus industry changing
- What we do know
  - If the change doesn't noticeably (10x rule of thumb) improve things it's not worth the change
- But what do we know really?
  - The old will be new again as folks reinvent or try previously tried approaches to new or even similar problems.



# The Future of RAC



- We can not, should not, and are not interested in competing with AWS, Azure, or GCP.
- Is there still a place for this when the big 3 have free tiers?
- We don't know for sure - but we do know for static workloads, self hosting outweighs cloud costs.

# Thank you!

Questions? Contact [rac-admin@cybera.ca](mailto:rac-admin@cybera.ca)

