The Distributed Cloud

Rick McGeer

University of Victoria

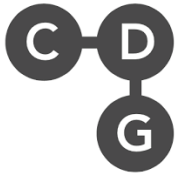

Takeaway

- The Distributed, Ubiquitous Cloud is the Internet of the Future
  - ***The most effective way to use a network is to send a program over it.***
  - Solution to the Zettaflood (soon to be the Yottaflood)
  - Enables big, interactive applications on small devices
  - New challenges: placement, security, context…

University of Victoria

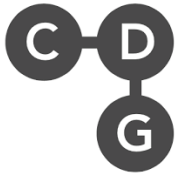

Outline

- **The Zettaflood**
- Big Data and Small Devices
- The Pollution Visualizer: A Case Study
- DEMO
- What is the Distributed Cloud?
- Examples of the Distributed Cloud
- Opportunities and Challenges

University of Victoria

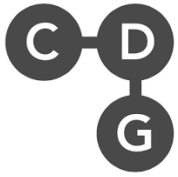

The Zettaflood

- A zettabyte added to the world's disks every 2 years
  - 10^{21} bytes, 10^{22} bits, 7 x 10^{7} seconds
  - 10^{14} bits/second, or 100 Terabits/second
  - *Just for ingress: egress is 10 x – 100 x ingress*
- We don't have the network to do that
  - Major NSF announcement: 100 Gb/s network for California Universities
  - *Can take < 0.1% of ingress data*

University of Victoria

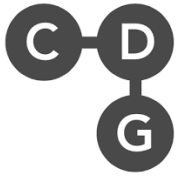

All Bandwidth is Local

- Typical Building:
  - outbound connection sized for 1 Mb/s/person.
  - Edge switch: 1 Gb/s/person
  - 99.9% of bandwidth internal
- City with FTTH:
  - 1 Gb/s/home, 50K homes: 50 Tb/s internal bandwidth
  - Intercity bandwidth *maybe* 100 Gb/s
  - 98%+ of network traffic is intracity

University of Victoria

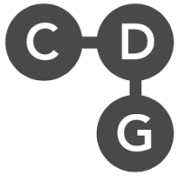

High Bandwidth Sensors Flood the Network

- iSight camera: 8 mp
  - 8 mp x 24 bits/pixel x 60 frames/sec = 11.4 Gb/s
  - 1.8 *billion* cameras in the world…
  - Cameras can generate 200 Exabits/second
- Many more examples….(VR, Internet of Things…)
- But…programs *reduce* data from sensors, *generate* it for end users
- Programs *much* smaller than the data they reduce or generate
- Bottom Line: Better to ship programs to data

University of Victoria

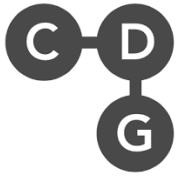

Outline

- The Zettaflood
- **Big Data and Small Devices**
- The Pollution Visualizer: A Case Study
- DEMO
- What is the Distributed Cloud?
- Examples of the Distributed Cloud
- Opportunities and Challenges

University of Victoria

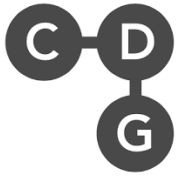

Big Data, Small Devices

- Want to render/visualize Big Data on tablets, cell phones, netbooks…
- Data/application must be resident in Cloud
- But “Classic Cloud” is too far away from the user
  - Amazon has 5 POPs in North America
  - 20ms – 50 ms away from most users
  - Need 1 ms – 10 ms for interactive applications

University of Victoria

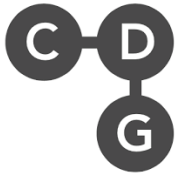

Why Cloud-Hosted Applications?

- Universal access on many clients
  - Handheld, laptop, netbook, VR headset, AR devices
- Control over execution environment
- Enhanced security for both application and user

University of Victoria

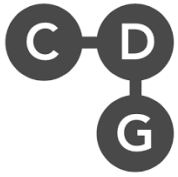

Outline

- The Zettaflood
- Big Data and Small Devices
- **The Pollution Visualizer: A Case Study**
- DEMO
- What is the Distributed Cloud?
- Examples of the Distributed Cloud
- Opportunities and Challenges

University of Victoria

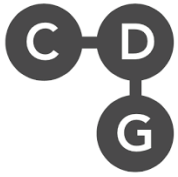

Visualize Big Data

Any Device

Fast

Anywhere

Collaboratively

Today: Pick 3

Distributed Cloud: All 5

University of Victoria

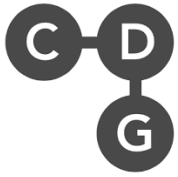

--

Distributed Collaboration Around Big Data

"A World Where Distance is Eliminated"
Experts Around the World Interacting With Data Visualizations as Easily as if They Were in the Same Room
Previously: Expensive Hardware (OptIPortal, CAVE) over Expensive, Special-Purpose Networks
Distributed Cloud: Any Device, Anywhere, Anytime, through a Web browser

University of Victoria

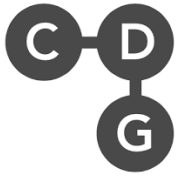

--

Collaboration Around Big Data

All About Size and Speed

- Data Set: 4 million points per month
- 100 MB/month

Too Much Data For Laptop

Way Too Much for Tablet/Phone/Netbook

Need Server Close to User

How Close Depends on Bandwidth

University of Victoria

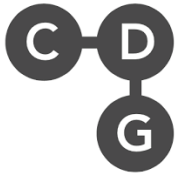

Size and Speed

Task: Draw 30,000 circles in 160 milliseconds

- World at 100-km resolution
- Quarter-continent at 10-km resolution

Requirement: 160-milliseconds

- User studies show dropoff beyond that.

Question. Can we do that from:

- Server on campus, Server in city, Server on continent, Single Server for World?

University of Victoria

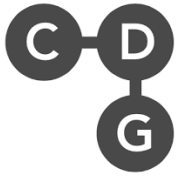

Size And Speed

	Server in Building	Server in City	Server on Continent	Worldwide Server
Request Time	1	5	50	250
Fetch Time	20	20	20	20
Transmit Time	8	30	300	1500
Draw Time	100	100	100	100
Total	129	155	470	1870

Time To Draw 30,000 Points in milliseconds. Goal: 150 ms

University of Victoria

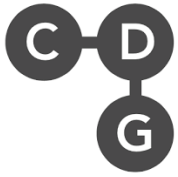

Size and Speed

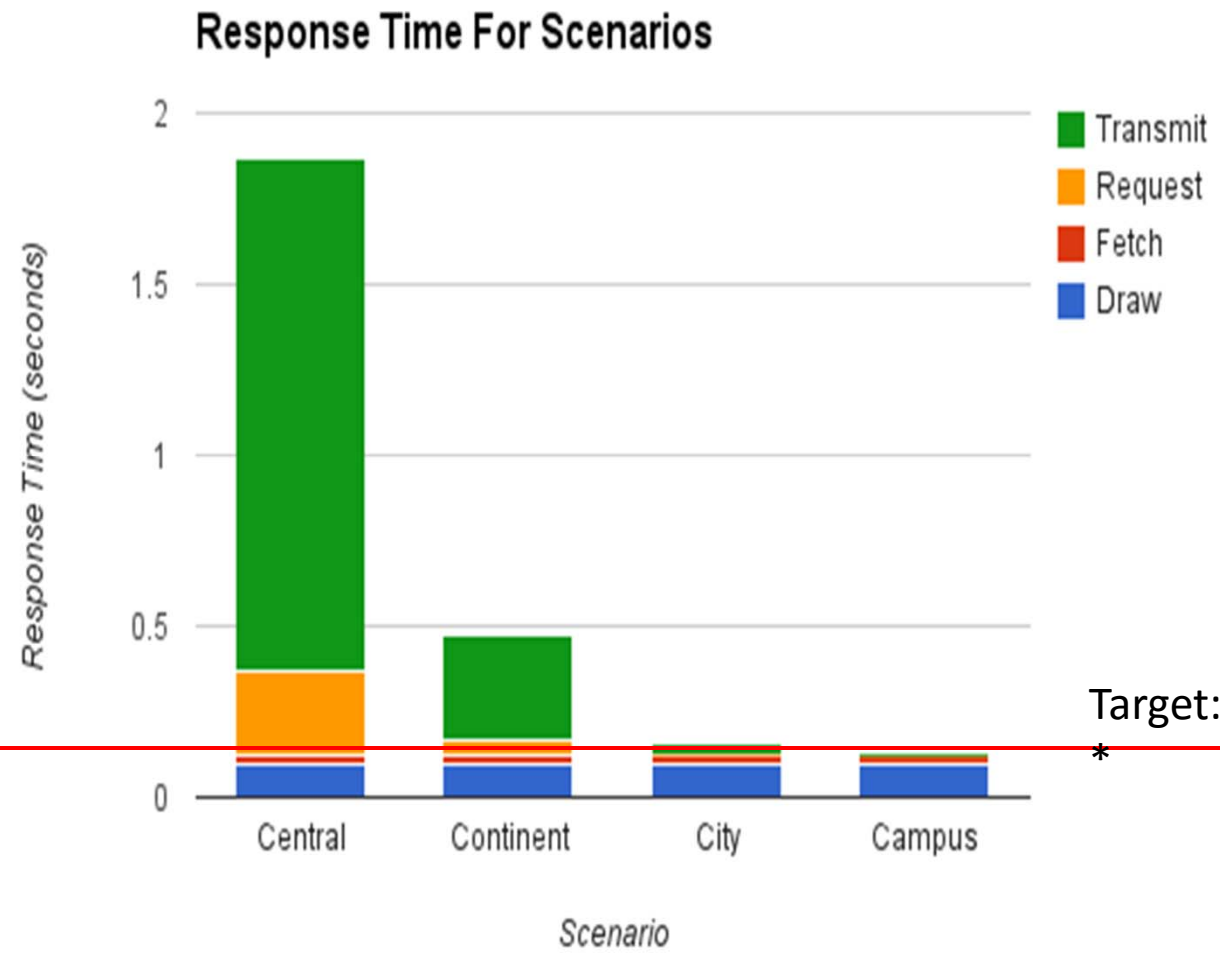

Un
of

nite

*N. Tolia, D. G. Andersen, and M. Satyanarayanan. Quantifying interactive user experience on thin clients. Computer, 39(3):46–52, 2006.

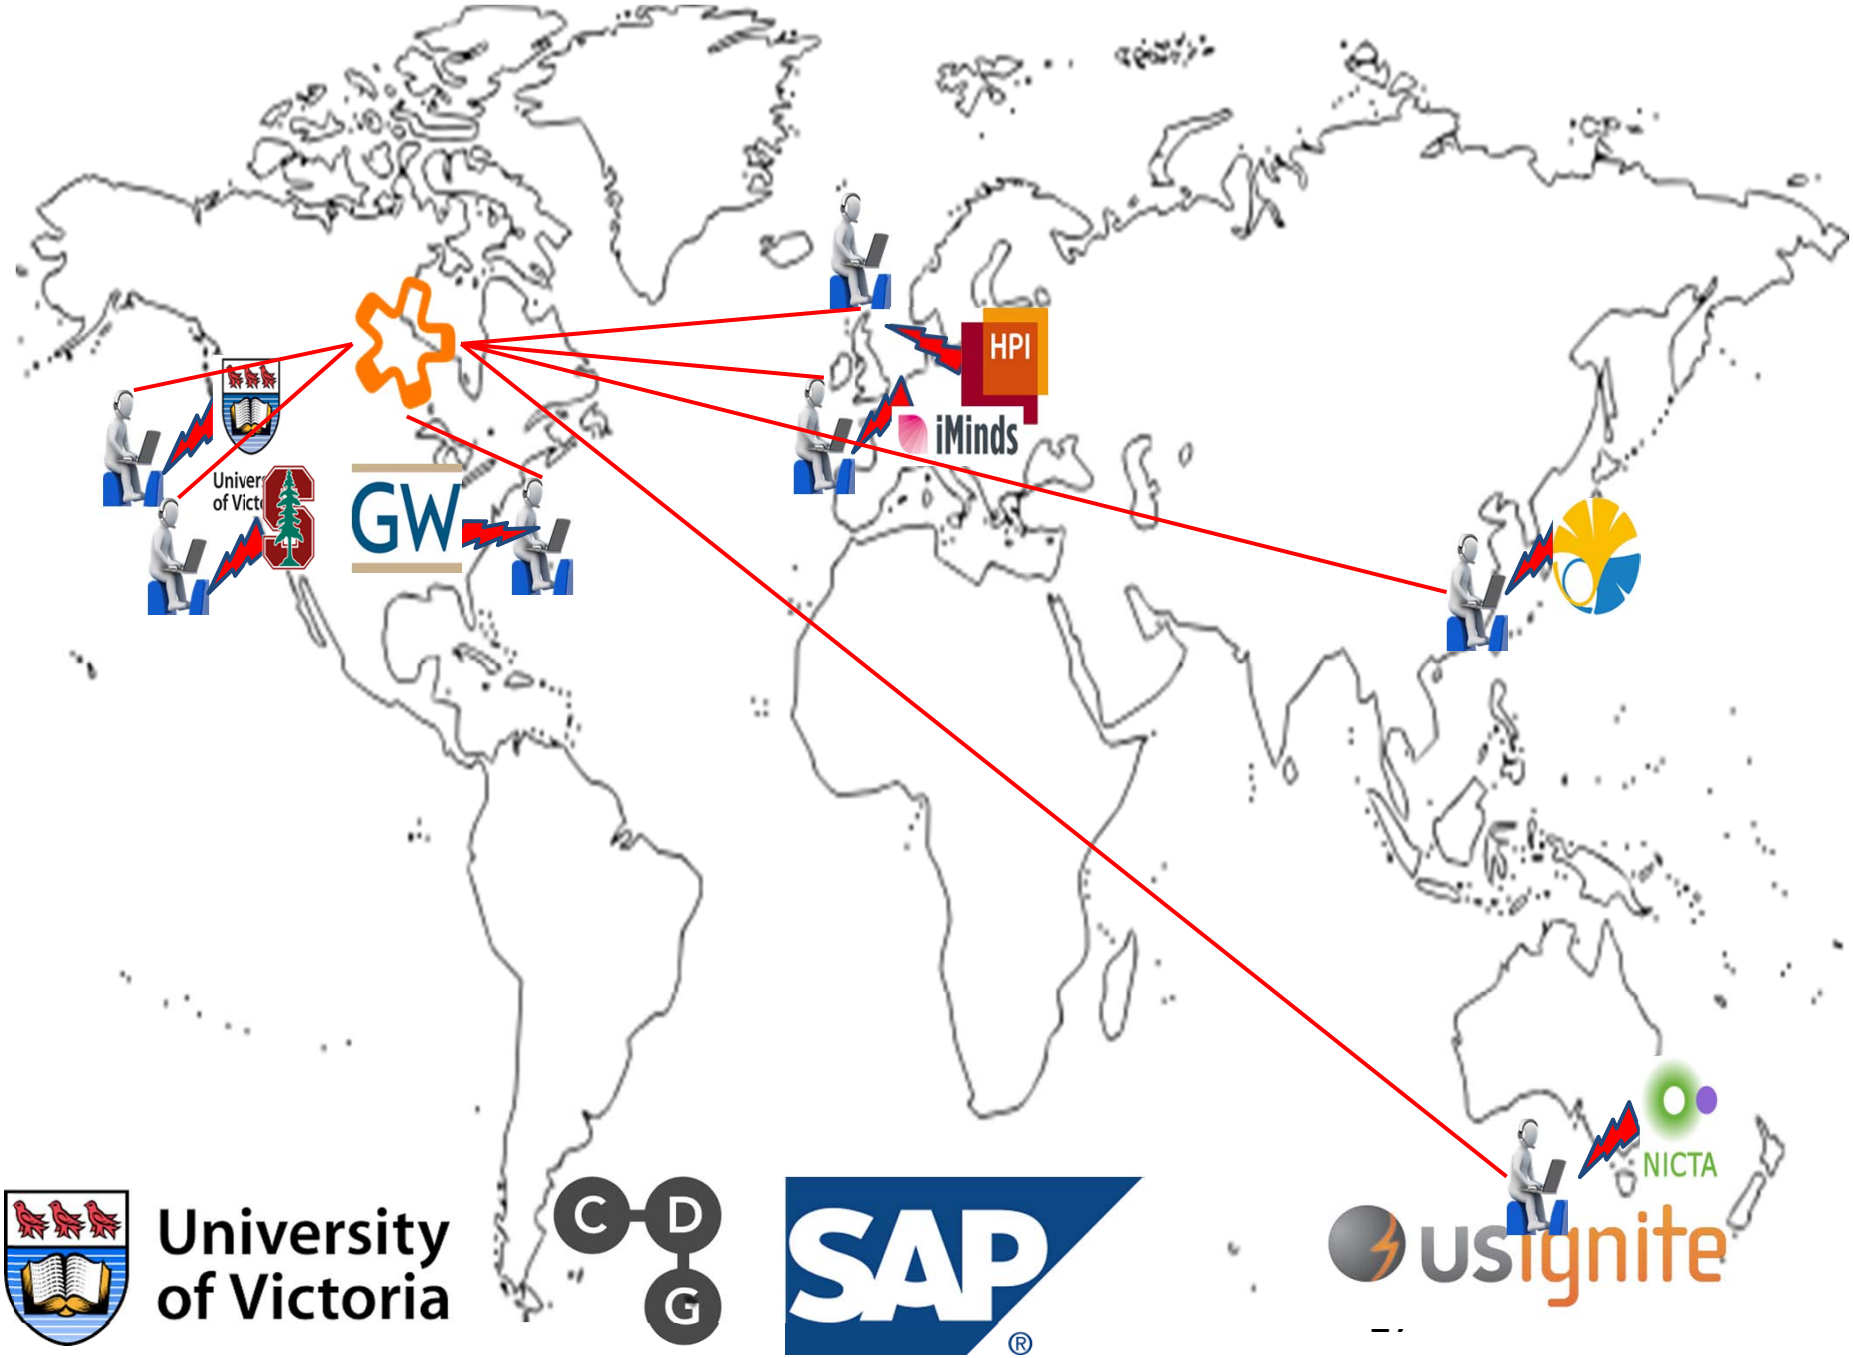
HPI
iMinds
GW
NICTA
University of Victoria
SAP
usignite

Outline

- The Zettaflood
- Big Data and Small Devices
- The Pollution Visualizer: A Case Study
- **DEMO**
- What is the Distributed Cloud?
- Examples of the Distributed Cloud
- Opportunities and Challenges

University of Victoria

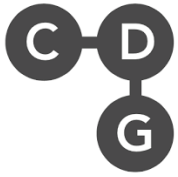

Outline

- The Zettaflood
- Big Data and Small Devices
- The Pollution Visualizer: A Case Study
- DEMO
- **What is the Distributed Cloud?**
- Examples of the Distributed Cloud
- Opportunities and Challenges

University of Victoria

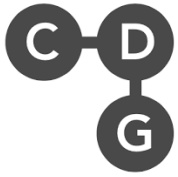

What Is The Distributed Cloud?

- First Approximation: EC2 with 50 PoPs across the US
  - Key difference: instantiate VMs in specific places, not just number of VMs
- Second Approximation: Layered Services

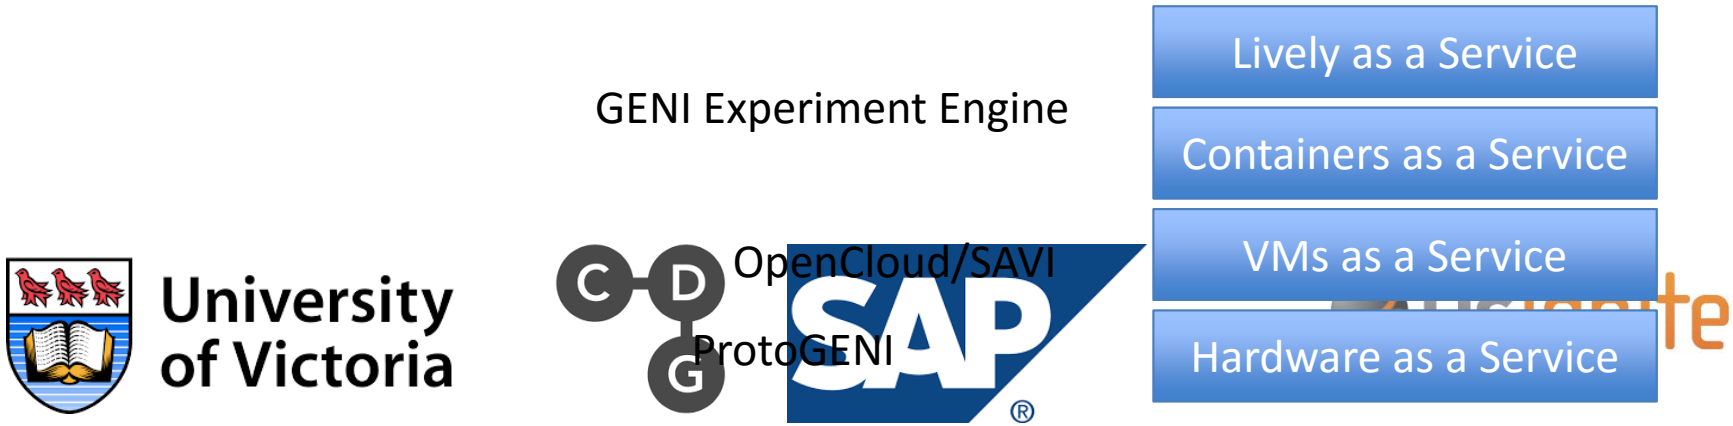

Outline

- The Zettaflood
- Big Data and Small Devices
- The Pollution Visualizer: A Case Study
- DEMO
- What is the Distributed Cloud?
- **Examples of the Distributed Cloud**
- Opportunities and Challenges

University of Victoria

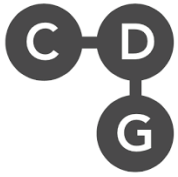

--

PlanetLab

- 1300 nodes at approx 400 sites worldwide
- Bare containers as a Service

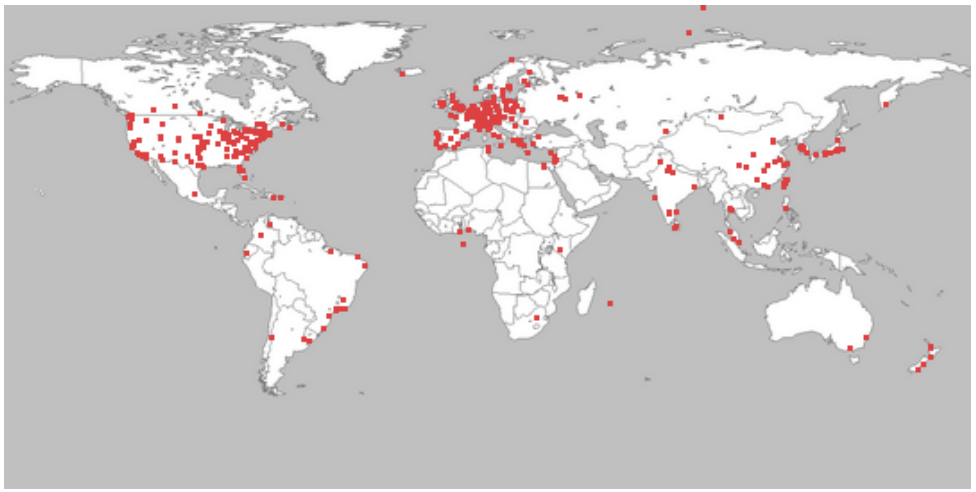

University of Victoria

C D G

SAP

usignite

GENI

- 50 baby racks across the US
- Full stack through ProtoGENI and the GENI Experiment Engine

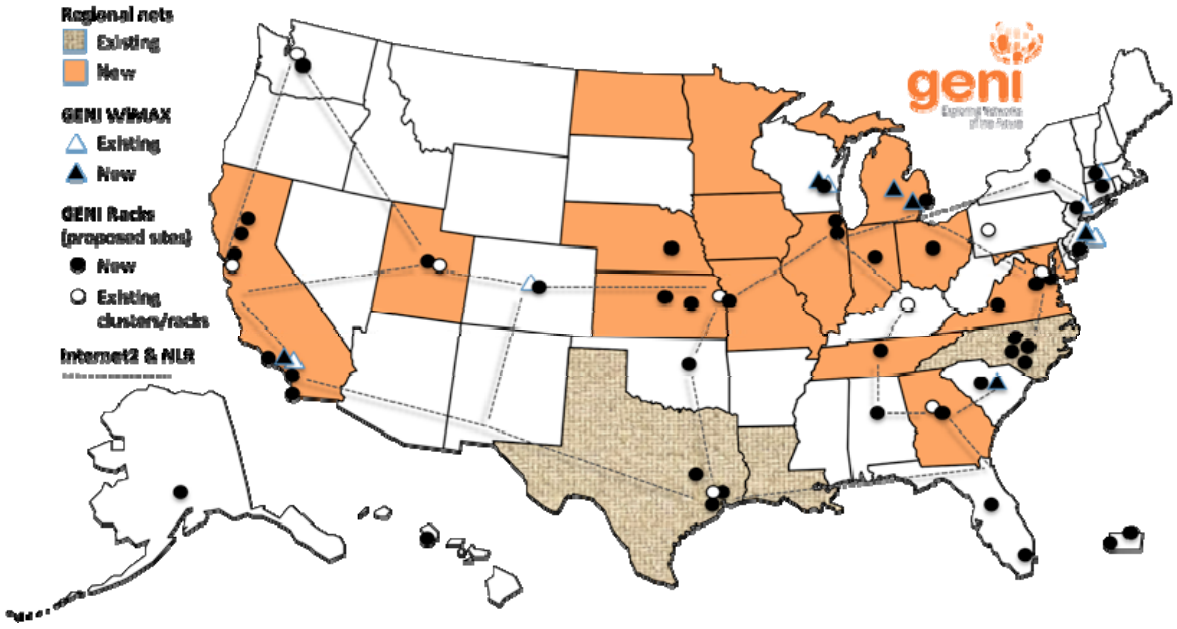

University of Victoria

usignite

SAVI (Canada)

- One “Core Node” (Large Rack)
- Several “Edge Nodes” (Baby Rack)
- OpenStack-based

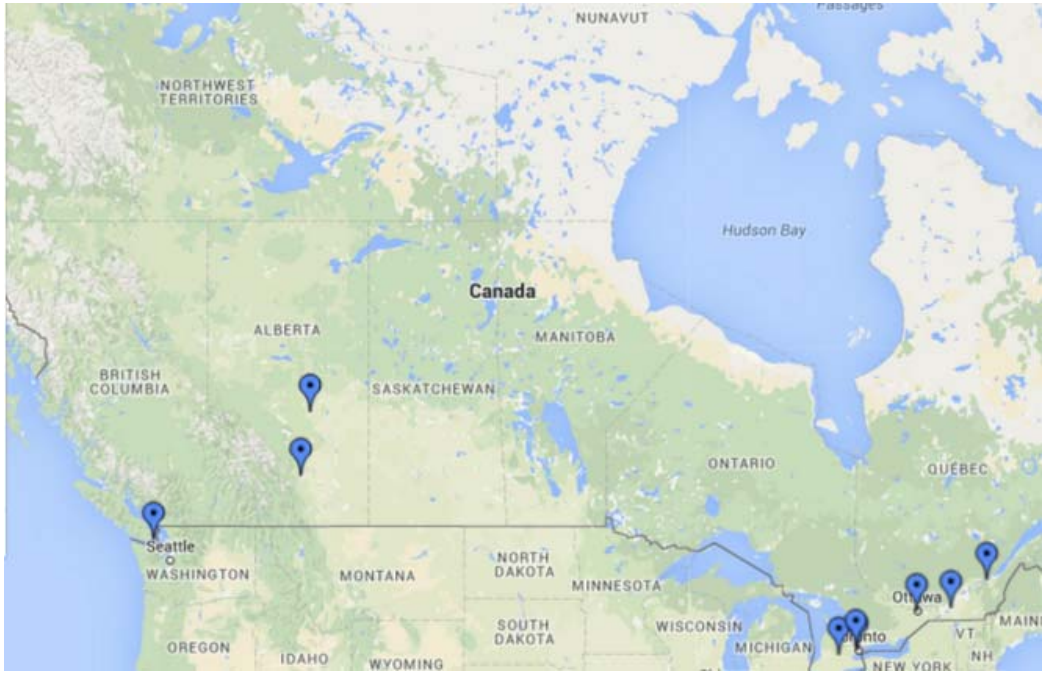

University of Victoria

C D G

SAP

usignite

FED4FIRE (EU)

- Federation of Cloud/Testbeds
- Based on GENI Software Stack

University of Victoria

C D G

SAP

usignite

Vnode/Flare (Japan)

As of 2/5/2015

Outline

- The Zettaflood
- Big Data and Small Devices
- The Pollution Visualizer: A Case Study
- DEMO
- What is the Distributed Cloud?
- Examples of the Distributed Cloud
- **Opportunities and Challenges**

University of Victoria

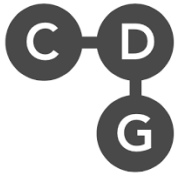

New Opportunities

- Collaboration Across Continents
- Application- and group-specific virtual internets
- Secure online education platform with real-time response
- Prototype for SmartCities/IoT
- Key technology behind 5G Wireless
  - 5G Wireless moves computation to the edge

University of Victoria

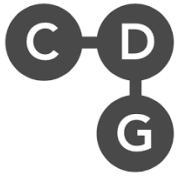

Key Problems To Solve

- Heterogenous ownership/administration
  - Key early experiment: federating SAVI and GENI
  - GENI users can now use SAVI testbed and vice-versa
  - Can build services spanning both infrastructures
- Heterogeneity key since it permits grassroots infrastructure
- Location- and context-aware programs
  - Know where and which infrastructure is being used

University of Victoria

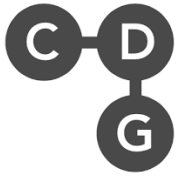

Takeaway

- The Distributed, Ubiquitous Cloud is the Internet of the Future
  - ***The most effective way to use a network is to send a program over it.***
  - Solution to the Zettaflood (soon to be the Yottaflood)
  - Enables big, interactive applications on small devices
  - New challenges: placement, security, context…

University of Victoria