Supporting Hybrid Services at an International Exchange Facility: *The Experience of Pacific Wave*

Dave McGaugh,

Pacific Northwest Gigapop

NANOG 39

Toronto, Canada

PACIFIC NORTHWEST GIGAPOP

About Pacific Wave

- A joint project between Corporation for Education Network Initiatives in California (CENIC), Pacific Northwest Gigapop (PNWGP), in cooperation with University of Southern California (USC) and University of Washington (UW)
- A distributed Internet Exchange Point (IXP) running the length of the entire United States Pacific Coast
- Supports high-end networking and protocols, including IPv4 (ucast/mcast), IPv6 (ucast/mcast), Jumbo Frames

Pacific Wave Initial Services

- Provide an open infrastructure for connecting international IP networks
- Facilitate any-to-any connectivity between connectors without the need for involvement of exchange operators
- Primary connectivity provided via shared VLANs
- Private VLANs used sparingly on as-needed basis

International Participation

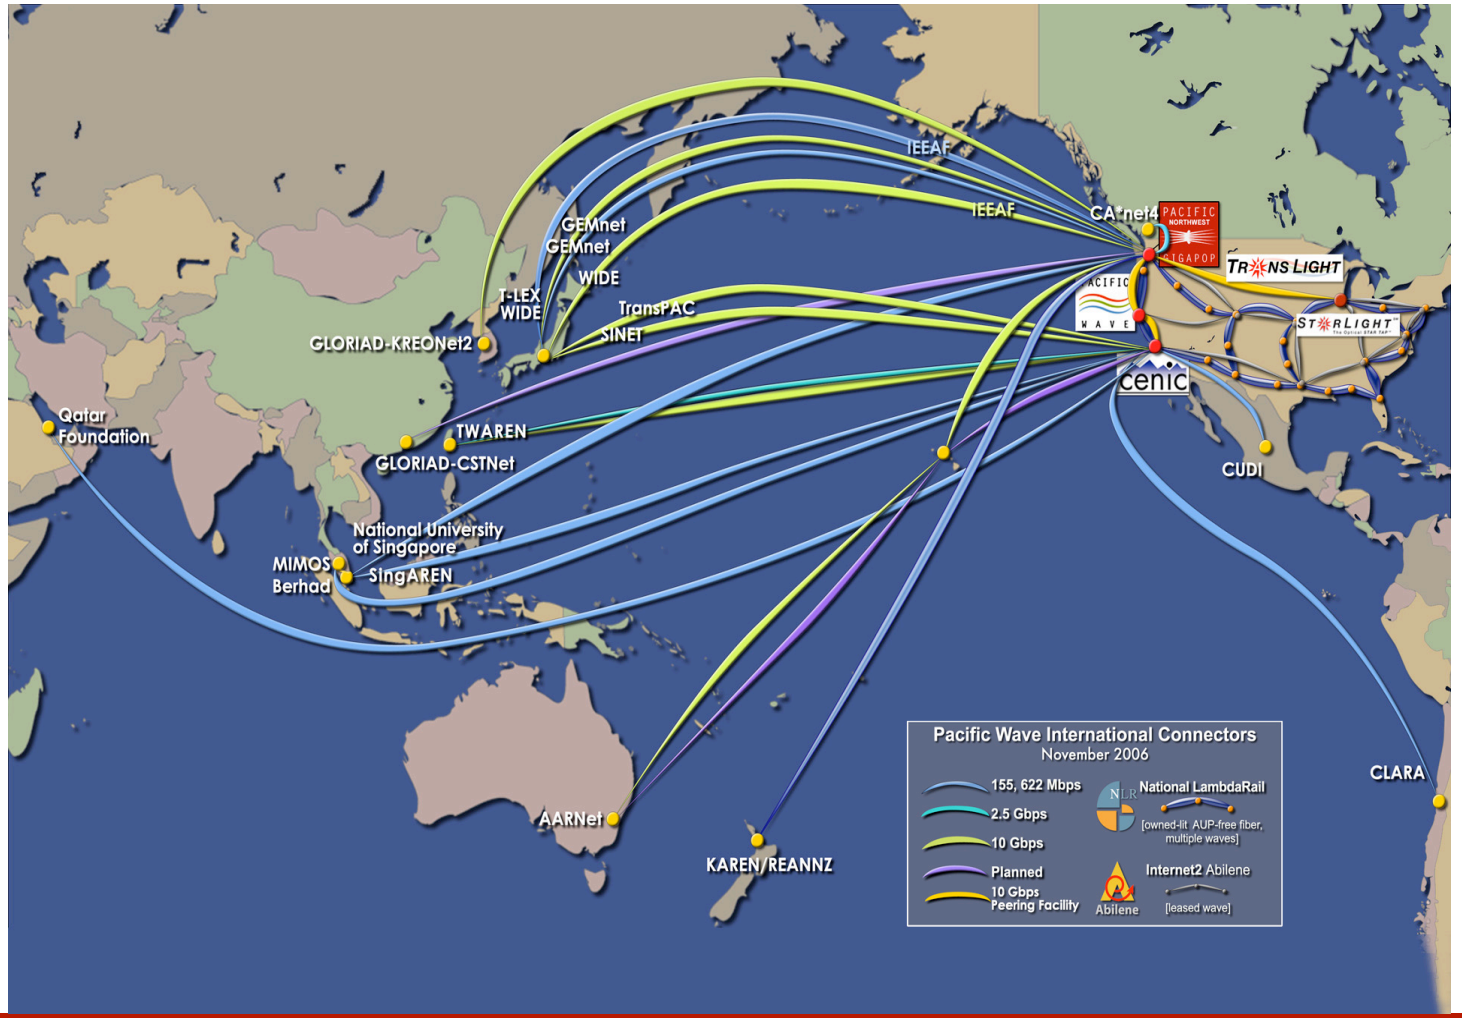

Characteristics of the Service Offering

- Distributed connectivity using shared broadcast domains
- Strict connector policy to ensure high availability and uptime
  - Single MAC per VLAN
  - No direct connection of layer-2 switching devices
  - Spanning tree root guard

Shared VLAN Infrastructure

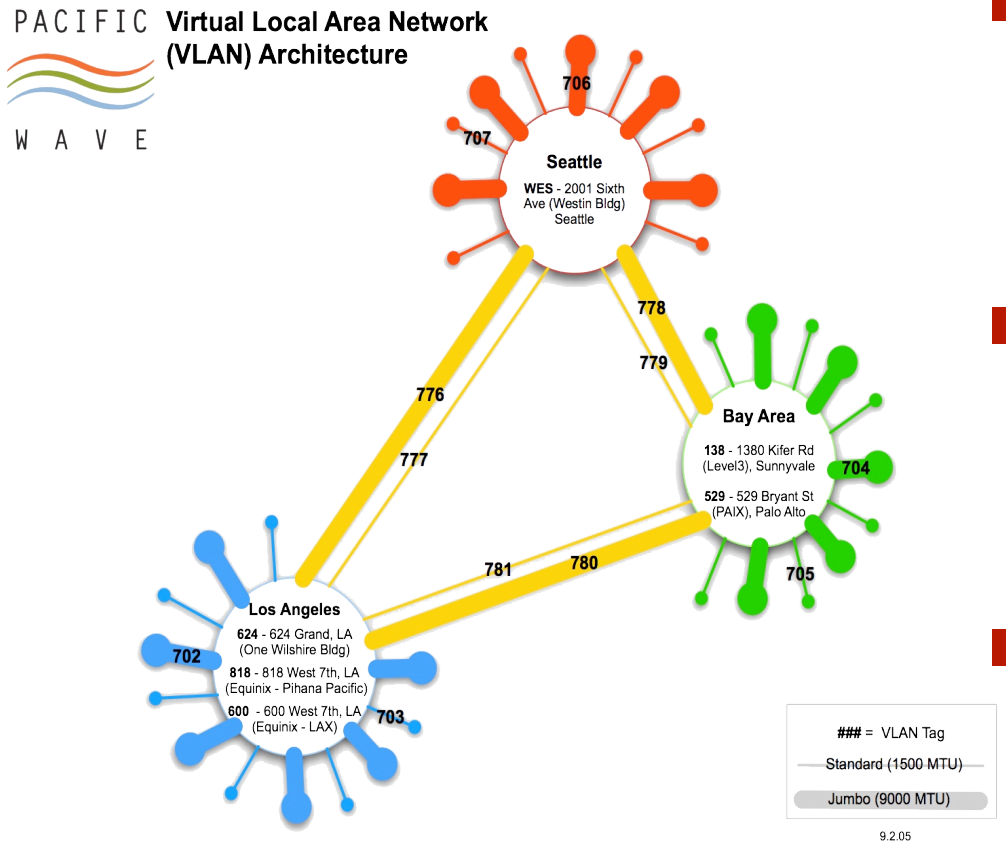

- Local, shared VLANs
  - Seattle
  - SF Bay Area
  - Los Angeles
- Intersite, shared VLANs
  - Seattle - SF Bay Area
  - SF Bay Area - LA
  - Seattle - LA
- Standard MTU and Jumbo MTU overlay

Physical Topology

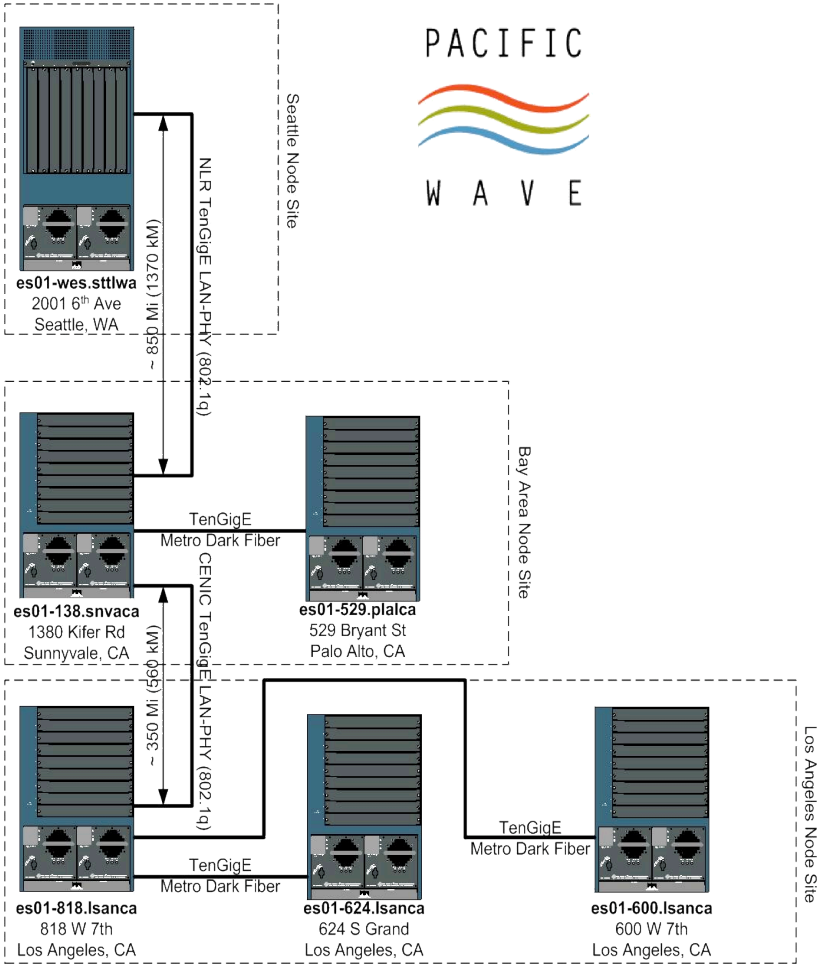

- Five 65/7600 layer-2 switches
- Deployed over NLR and CENIC 10g lambdas
- Seattle
  - Westin
- SF Bay Area
  - Sunnyvale (Level 3)
  - Palo Alto (PAIX)
- Los Angeles
  - One Wilshire
  - Equinix (600 W 7th)
  - Equinix (818 W 7th)

First Major Use of Hybrid Exchange Services

- First wide use of hybrid services: iGrid 2005
  - Over 10 Gbps bidirectional traffic (3 Gbps multicast) coexisted with production exchange traffic without detriment
  - Made use of both shared VLAN and private VLAN service providing “lightpath” extension
  - Several short-lived private VLANs used and immediately torn down post-conference
- This demonstrated the value in using Pacific Wave layer-2 switching devices, previously reserved for layer-3 peering only as a lightpath grooming facility

Lightpath Definition

- A lightpath is a point-to-point channel which the user takes as a dedicated line with a given service level guarantee.
  - SONET/SDH TDM channel
  - WDM wavelength
  - MPLS layer-2 VPN
  - Any concatenation of the above
- Only one hop between the endpoints of a lightpath!

iGrid 2005 Network Augment

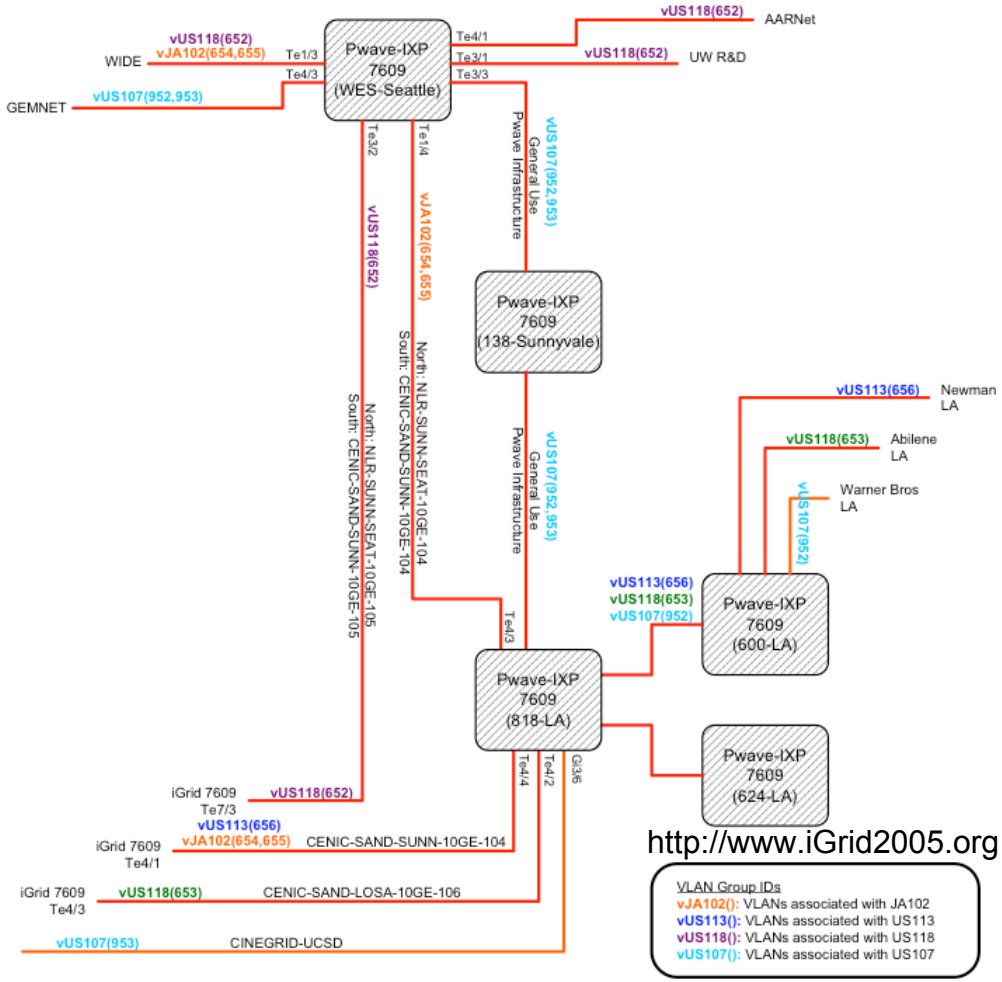

http://www.iGrid2005.org

- Two additional 10g lambdas provisioned via NLR and CENIC
- International connectivity made via Seattle node
- Additional lambdas in Southern California provisioned to serve the iGrid venue and local demo connectivity

The Administrative Leap to a Hybrid Exchange

- Connector policy must be relaxed to allow lightpaths
  - How can this be done without undermining the reason this policy exists to begin with?
- Need to react more quickly to provisioning requests
- Should new billing models be considered?

What This Means Technically

- Increased use of Private VLANs
- Less control over what type of devices connect to the equipment, both locally and across the wide-area
- Requires much more coordination when establishing VLANs to avoid VLAN ID collision when interconnecting multiple layer-2 domains

The Reality

- In many cases, lightpaths are not provided over a single technology end-to-end
- The exchange facility can fill the role of this technology translation point with a robust toolkit of layer-1 and layer-2 devices aside from just an any-to-any meet point of networks

Layer-1 Infrastructure

- Seattle
  - Nortel HDXc co-operated with Canarie
  - Cisco IEEAF ONS 15454 MSPP
    - 4 ML1000-2 line cards (Capable of GFP-F and LEX)
  - KREONet2 ONS 15454 MSPP
- Los Angeles
  - CENIC ONS 15454 MSPP
  - AARNet ONS 15454 MSPP

Combined Layer-1 and Layer-2 Infrastructure

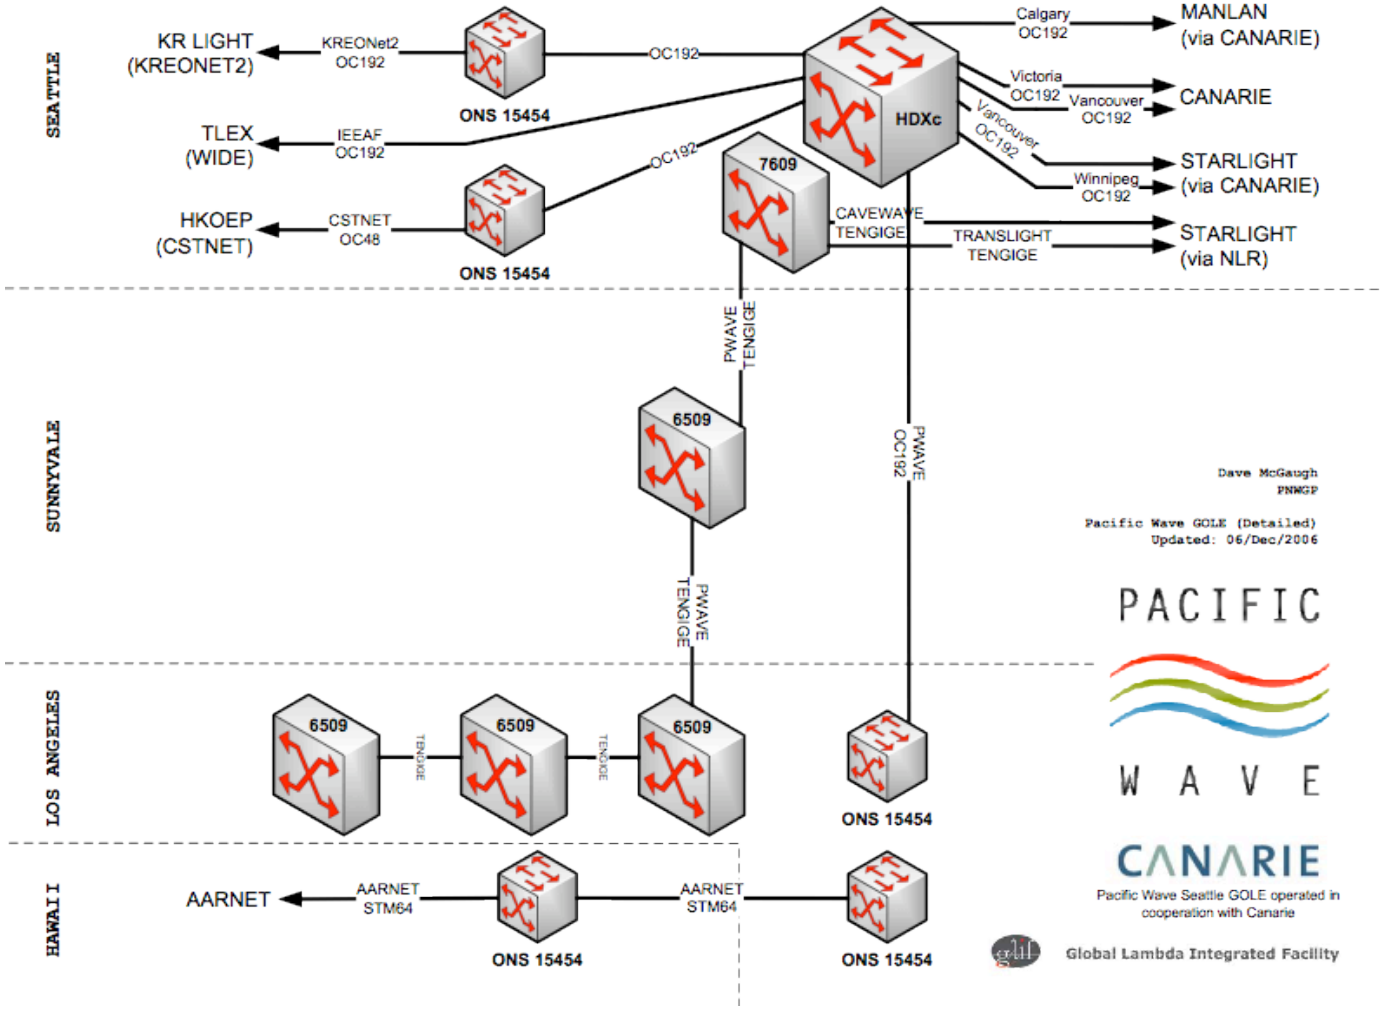

Supported Lightpath Types

- Ethernet frame-mapped onto Layer-1
  - Provides no stat-muxing functionality
  - Reconfiguration causes a circuit outage
  - Predictable Jitter and Latency
- Switched Ethernet
  - Ability for TenGigE convenience interfaces when 10 Gbps of bandwidth is not available/required
  - May introduce frame loss on congested links or backplanes
  - Some debate as to whether this fits the true definition of “lightpath”

The Hybrid Approach

- The hybrid exchange operator can provide easy interconnection between Layer-1 and Layer-2 devices as necessary
  - Today: manual with physical cross-connect
  - Future: optically switched manually or via control plane protocols
- Technology choices made based on:
  - Performance requirements
  - Bandwidth requirements
  - Timeframe desired for lightpath activation
  - Availability of wide-area transport

Pacific Wave's Involvement in GLIF

- GLIF is the Global Lambda Integrated Facility
  - An international virtual organization that promotes the paradigm of lambda networking
- GOLE is a GLIF Open Lightpath Exchange, some others include:
  - KRLight in Daejeon
  - NetherLight in Amsterdam
  - HKOEP in Hong Kong

A Distributed GOLE?

- A term as yet formally defined in the GLIF community
- Properties of a distributed GOLE
  - Multi-transport interconnect between nodes
    - Layer-1 lambda
    - Layer-2 frame mapped (i.e. GFP-F)
    - Layer-2 switched
  - Able to establish lightpaths between and through the nodes with the efficiency of a single operating entity
  - No bandwidth limitations between nodes (within reason…)

Pacific Wave GOLE Diagram

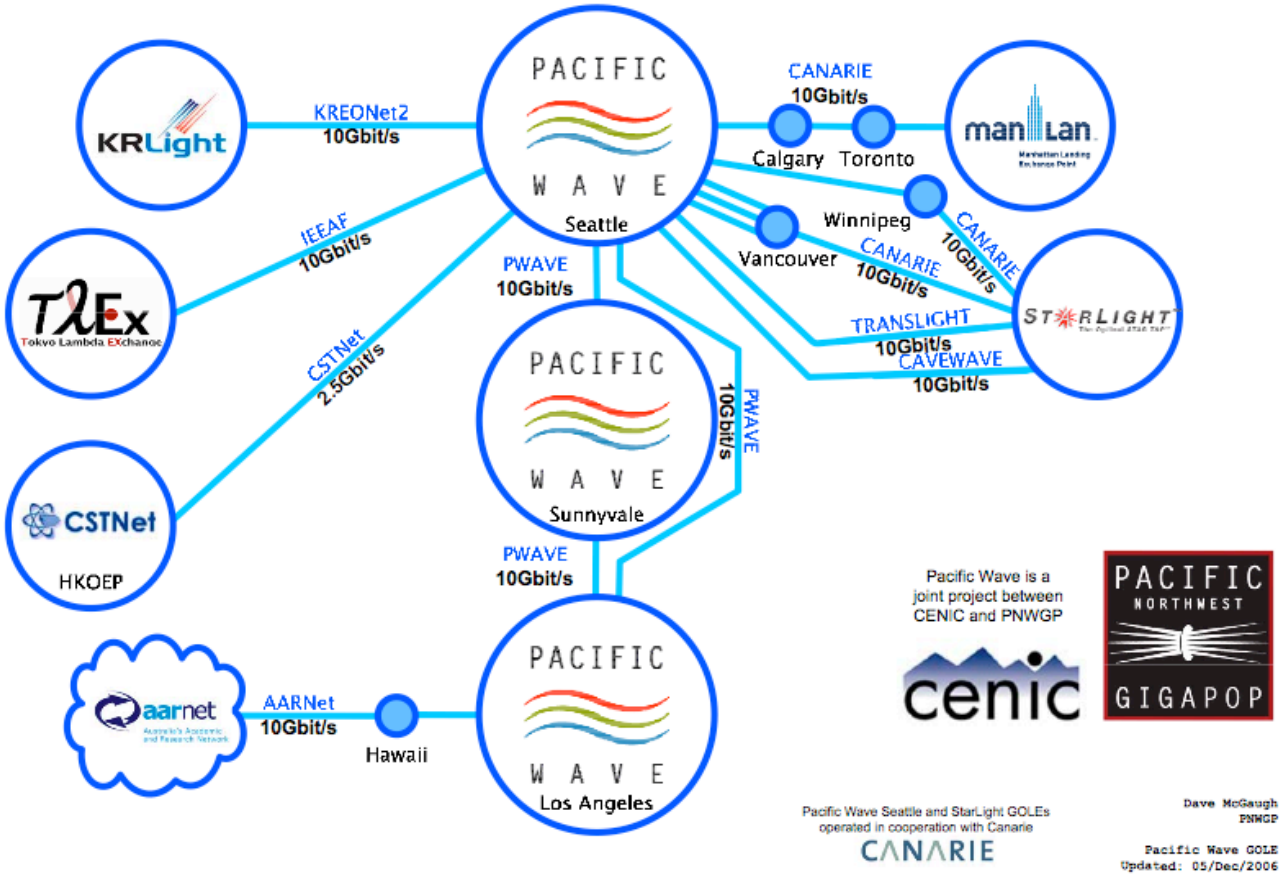

Future

- Today, Pacific Wave Layer-2 switching is a series of linear interconnected fabrics
- Perhaps more complex wide-area topologies may one day be necessary on the Layer-2 network
  - Bandwidth management
  - Resiliency
- MPLS layer-2 Ethernet VPNs may provide additional opportunities to support future requirements
- Optical switching facilitates efficient lightpath establishment, especially during technology translation, i.e. frame-mapped ⇒ switched Ethernet

Upcoming Technologies

- GFP-F TenGigE LAN-PHY using VCAT/LCAS
  - Might negate some of the need for Ethernet switching, such that TenGigE convenience interfaces can be mapped to sub-10 Gbps SONET/SDH circuits
- More robust VLAN tag rewrite capabilities
  - VLAN ID space management is becoming more and more challenging as Layer-2 facilities continue to interconnect
  - Some of these challenges can be solved with MPLS L2VPNs

More Information

- Pacific Wave:
  - http://www.pacificwave.net
- GLIF Website:
  - http://www.glif.is
  - http://wiki.glif.is
- Network Description Language
  - http://www.science.uva.nl/research/sne/ndl
- NLR (National LambdaRail)
  - http://www.nlr.net

PACIFIC

WAVE

Thank You!

dmcgaugh@pnw-gigapop.net

info@pacificwave.net