# UCB External Connectivity

ken lindahl
communication and network services
uc berkeley
lindahl@berkeley.edu

EE122, 7 May 2004

## CalREN-2: 1998 to 2003

- conceived as a state-wide high-performance, advanced services network serving California's research and higher education communities.
- charter members: UC, CSU, Stanford, CalTech, USC
- CalREN-2 network was supposed to OC-48 (2.5Gbps), but what got delivered was two OC-12 (622 Mbps) SONET rings
  - should have asked for "OC-12**c**." oops.
- implemented as one POS (Packet Over SONET) ring and one ATM ring



## CalREN-2 northern gigaPOP, ca. 11/98

- "distributed gigaPOP," with backbone equipment (SONET MUXes, ATM switches, IP routers) located in campus facilities.
- commercial SONET service for connections between campuses.

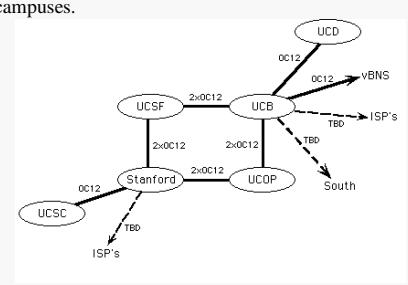




## CalREN-2 southern gigaPOP, ca. 11/98

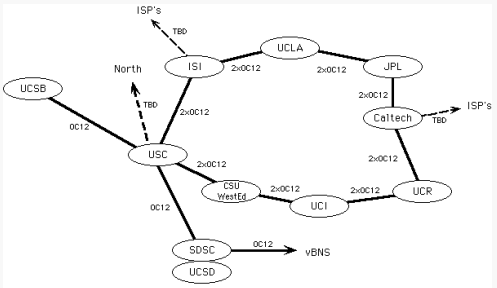


Figure 2: CalREN-2 LA Basin GigaPOP Bandwidth Plan



## CalREN-2 backbone, circa 11/98

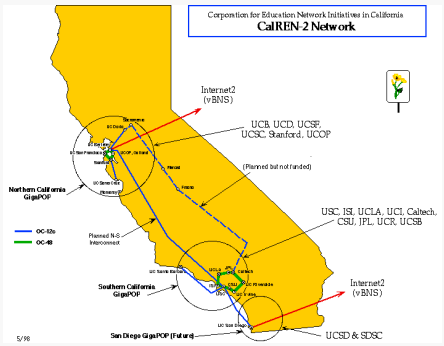




## CalREN-2: 1998 to 2003

- UCB initially connected to CalREN-2 with a single OC-12 POS link
- subsequently added one, then two, OC-3 (155 Mbps) ATM links for ISP traffic
  - separate path required in order to control campus ISP costs
  - diverse routing provided a measure of redundancy, for protection against failures (but Evans Hall Machine Room was still single point of failure)

## UCB external connectivity, ca. 5/2003 (1)

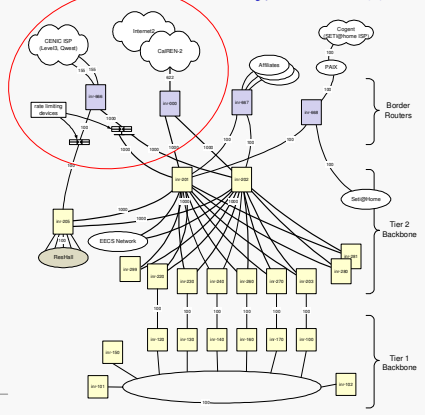




## shortcomings of the original CalREN-2 network

- campus facilities often not hardened against power outages, flooding, rats, … (read: Evans Hall)
- very little redundancy in network; failures tended to have wide-spread, significant effects
- campus support staff often not funded for 24x7 support
- commercial SONET circuits very expensive, especially after initial 3-year contract expired
- mixing research and ISP traffic on a single backbone had an unexpected result: reliability and availability (for ISP) became more important than flexibility (for research), it became very difficult to deploy new features on the backbone.



## CENIC next generation backbone considerations (1)

- running out of bandwidth on OC-12 backbone ciruits
- eliminate dependence on flakey campus facilities
- provide redundant connections to campuses
- need to separate production and research networks

| Production Networks | Advanced Services Networks |
|---|---|
| avoid change whenever possible | constantly changing |
| run well-tested code, let somebody else find the bugs | run new code as soon as it becomes available |
| very low tolerance for outages | outages are undesirable, but acceptable |
| capacity planning based on expected long-term utilization | capacity planning based on large amounts of headroom needed for occasional use |

- need for true separation of ISP traffic for rate-limiting



## CENIC next generation backbone considerations (2)

- over-investment in fiber plants causing many carriers to sell or lease fiber at low cost
  - vs. recurring high costs for commercial circuits
- development of wave division multiplexing technology
  - providing multiple high-speed (2.5 Gbps, 10 Gbps, 40 Gbps, …) channels over a single fiber pair (single fiber strand in some cases)
- researchers requesting "private pipes" between campuses
  - experiment with non-IP (unroutable) packets
  - experiment with routing protocols
  - experiment with "dangerous" packets
- small number of researchers requesting dark fiber between campuses



## the pyramid!

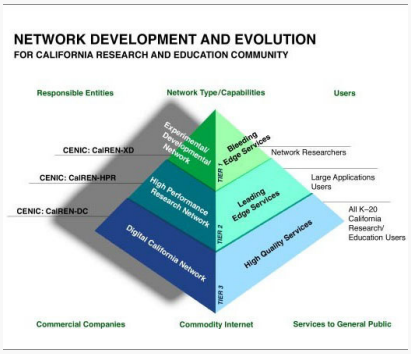


- plus, ISP "sub-layer" within CalREN-DC



## CalREN Optical Backbone, ca. 3/2004

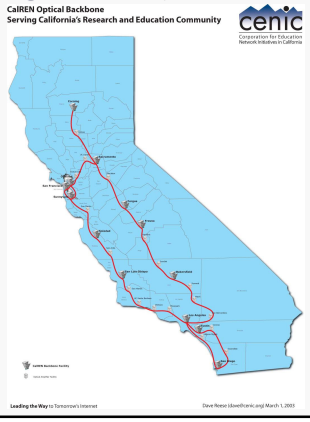

## Wave Division Multiplexing (sort of) demystified

- multiple channels of data carried over single fiber pair between WDM terminals
- each channel has its own wavelength ("lambda" or "wave")
- optical multiplexers combine multiple wavelengths into single composite signal; optical demultiplexers split composite signal into individual wavelengths
- amplification and regeneration needed for longer distances

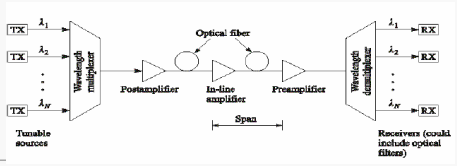




## CalREN backbone waves

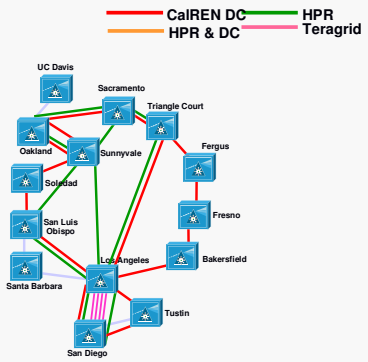




## Bay Area Metro Ring

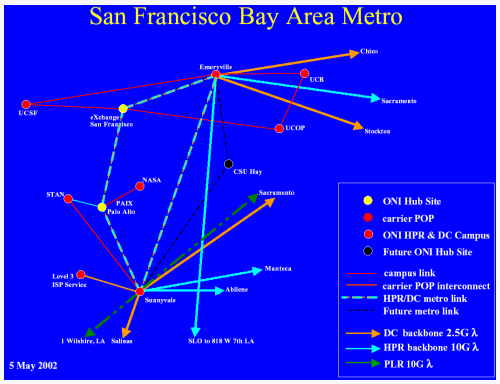




## CalREN-DC: 10/2003 to present

- OC-48c (2.5 Gbps) backbone over private DWDM fiber
- Six GE links to the commodity Internet (Qwest, Level3, Cogent)
- "no-cost" private peerings at PAIX, LAAP, SDNAP
- Serving:
  - California State University (all campuses)
  - California Community Colleges (all campuses)
  - University of California (all campuses)
  - ~80% of K-12 schools in California
- Estimated user base is 10 million faculty, staff and students!

  AOL: 26 million users
  **CalREN-DC: 10 million users**
  MSN: 9 million users
  Earthlink: 5 million users
  Netzero/Juno: 5 million users
  Comcast/ATT: 3.6 million users
  SBC Internet Services: 2.2 million users



## CalREN-HPR: 2/2004 to present

- 10Gigabit Ethernet backbone over private DWDM fiber
- Cisco 12400 routers
- 10GE connection to Abilene
- connection to CUDI (Mexico's version of Internet2)
- Jumbo frames (9180) and IPV6 enabled
- Serving:
  - University of California (all campuses)
  - Stanford
  - University of Southern California
  - CalTech & Jet Propulsion Laboratory (JPL)
- (future) 10GE to campuses, 12-18 months from now
- (future) PacificWave: distributed exchange point with POPs between LA and Seattle, ideal for landing trans-Pacific fiber links



## UCB connectivity to CalREN (1)

- Total of six GE links:
  - GE to CalREN-HPR in CEV near Recreational Sports Facility**
  - GE to CalREN-DC in CEV near Recreational Sports Facility
  - GE to CalREN-ISP in CEV near Recreational Sports Facility
  - GE to CalREN-HPR in Evans Hall**
  - GE to CalREN-DC in Evans Hall*
  - GE to CalREN-ISP in Evans Hall*

  * share a single 1 Gbps channel over SBC EGM
  ** not yet in service, expect service June 2004
- Redundancy: two campus locations, connected to different CalREN POPs via diverse fiber paths
  - OAK POP (Emeryville) via dark fiber leased from MFN
  - SFO POP via SBC Extended GigaMan (EGM) service

## new UCB border routers

- existing border routers were not up to the task
  - inr-666, a cisco 7507, can handle GE interfaces but with only 300-400 Mbps onto the backplane.
  - inr-000, a cisco GSR 12008, can handle multiple GE interfaces, but at great expense, and the 12008 is already being "sunsetted" by cisco.
- requirement for redundancy means at least two routers
- need to maintain separation of ISP from non-ISP traffic would have require (at least) four border routers, rather than just two…
- … or use "Policy Constrained Routing" developed in Internet2 Routing WG.
- RFP for new campus border routers led to selection and purchase of two Juniper M40e routers
  - being installed and configured now, in full production Summer 2004



## UCB connectivity to CalREN (2)

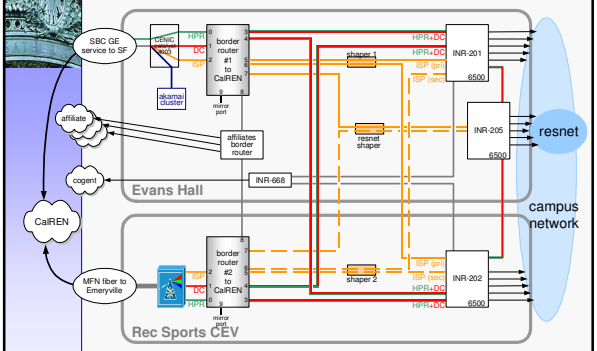


Figure 2. UC Berkeley Border Topology showing connections to next-generation CalREN network.



## Policy Constrained Routing

- developed by Cisco and Juniper, working with CENIC and Internet2 Routing Working Group to address limitations of traditional destination-based routing decisions
- uses "virtual routing and forwarding" tables (VRFs) to make forwarding decisions based on router input interface and destination IP address
  - a specific VRF is applied to all packets entering a particular interface
  - different VRFs, expressing different routing policies, applied to different interfaces
- needed at UCB to maintain separation of ISP vs non-ISP traffic



## Policy Constrained Routing in UCB border routers

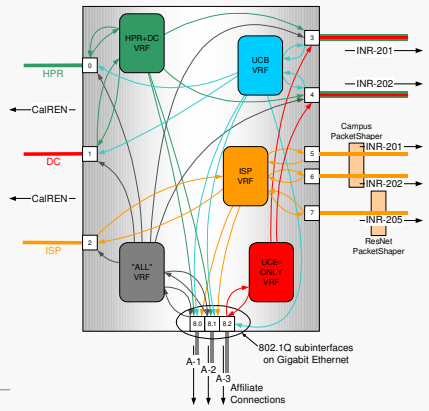




## National LamdaRail

- coalition of E&R entities, including CENIC, PNWGP, I2, Cisco
- uses DWDM (40 10GE lamdas) for multiple testbed networks
  - waves (L1) or switched ethernet (L2) or routed IP (L3
- first transcontinental Ethernet network
- additional dark fiber pairs