

ESnet On-demand Secure Circuits and Advance Reservation System (OSCARS)

Chin Guok
Network Engineering Group

Energy Sciences Network
Lawrence Berkeley National Laboratory

ESCC
July 23 2008

Networking for the Future of Science



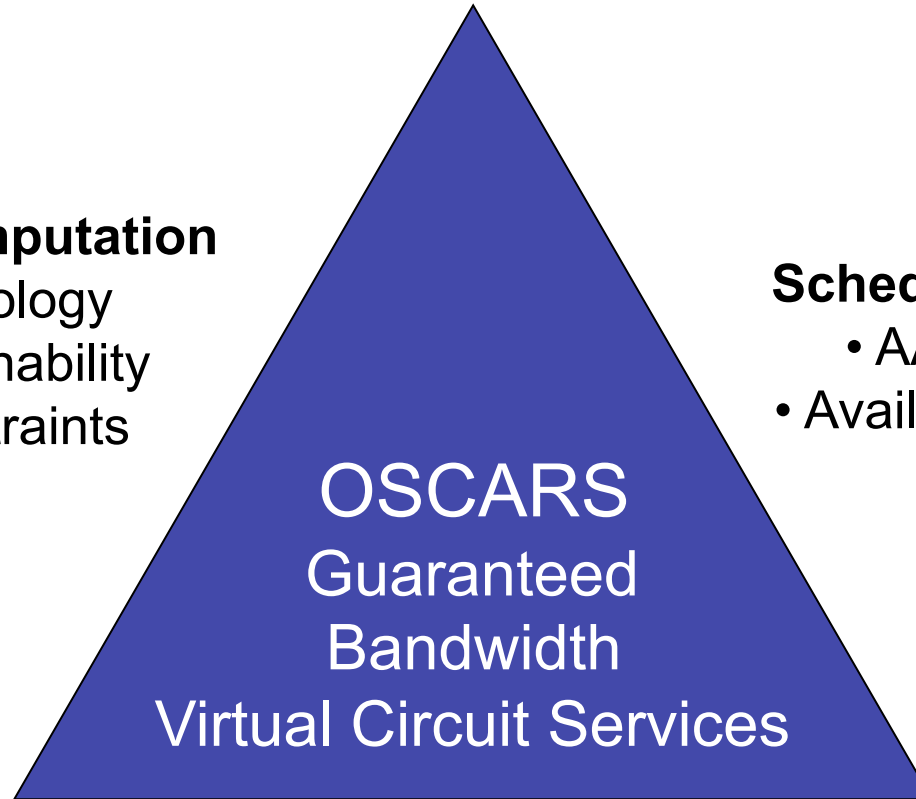
OSCARS Overview

Path Computation

- Topology
- Reachability
- Constraints

Scheduling

- AAA
- Availability



Provisioning

- Signaling
- Security
- Resiliency/Redundancy

Using OSCARS

- Web-Based User Interface (WBUI)
 - SSL connection to server
 - Username and password login

- SOAP Messages
 - SSL connection to server
 - WSDL service description
 - Signed SOAP messages

```
<?xml version="1.0" encoding="UTF-8"?>
<wsdl:definitions name="OSCARS"
  targetNamespace="http://oscars.es.net/OSCARS"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/"
  xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap12/"
  xmlns:tns="http://oscars.es.net/OSCARS">

  <wsdl:documentation>
    This is WSDL for the OSCARS public reservation interface. These
    messages must be signed using the following WS-security
    standards. The message is time stamped and includes the X.509
    certificate of the signing entity. The timestamp, certificate
    and message body are all signed. DRAFT V1.0 Nov 2006
  </wsdl:documentation>

  <!-- Element definitons -->
  <wsdl:types>
    <xsd:schema targetNamespace="http://oscars.es.net/OSCARS"
      elementFormDefault="qualified"
      xmlns:xsd="http://www.w3.org/2001/XMLSchema"
      xmlns:tns="http://oscars.es.net/OSCARS">
      <xsd:include schemaLocation="OSCARS.xsd" />
    </xsd:schema>
  </wsdl:types>

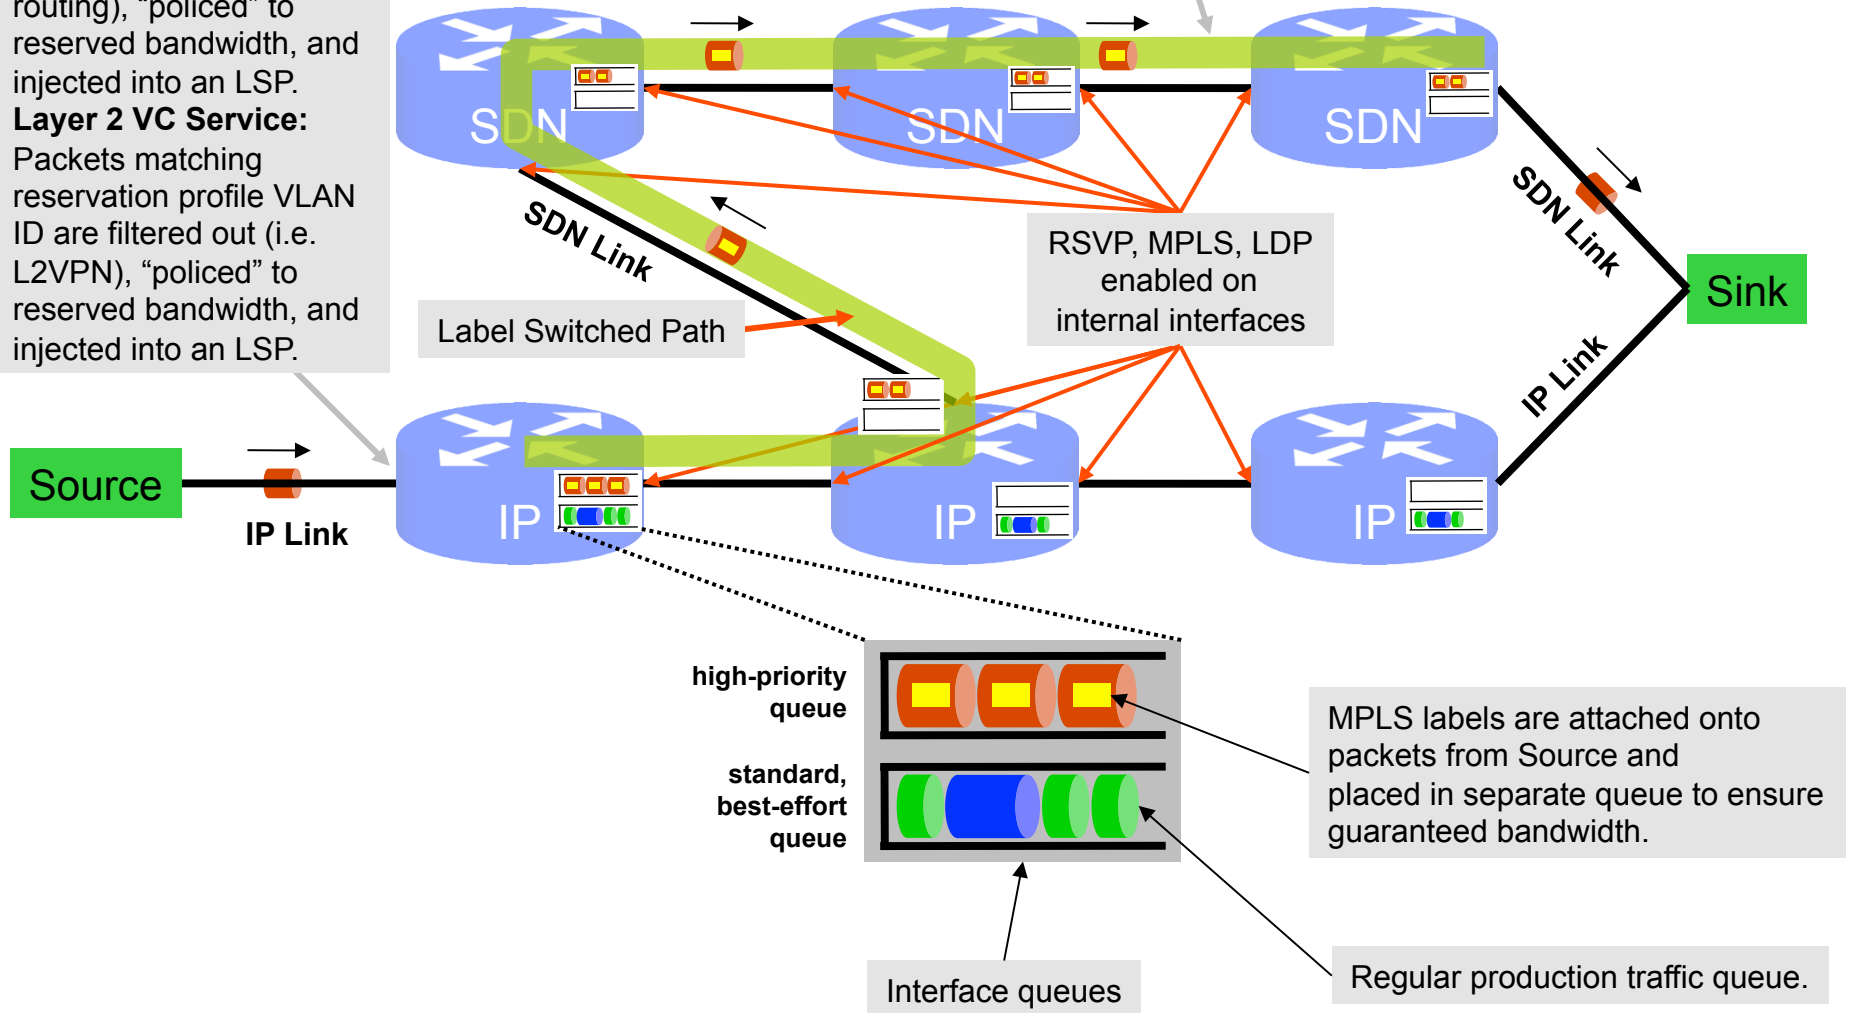
  ...
```

The Mechanisms Underlying OSCARS

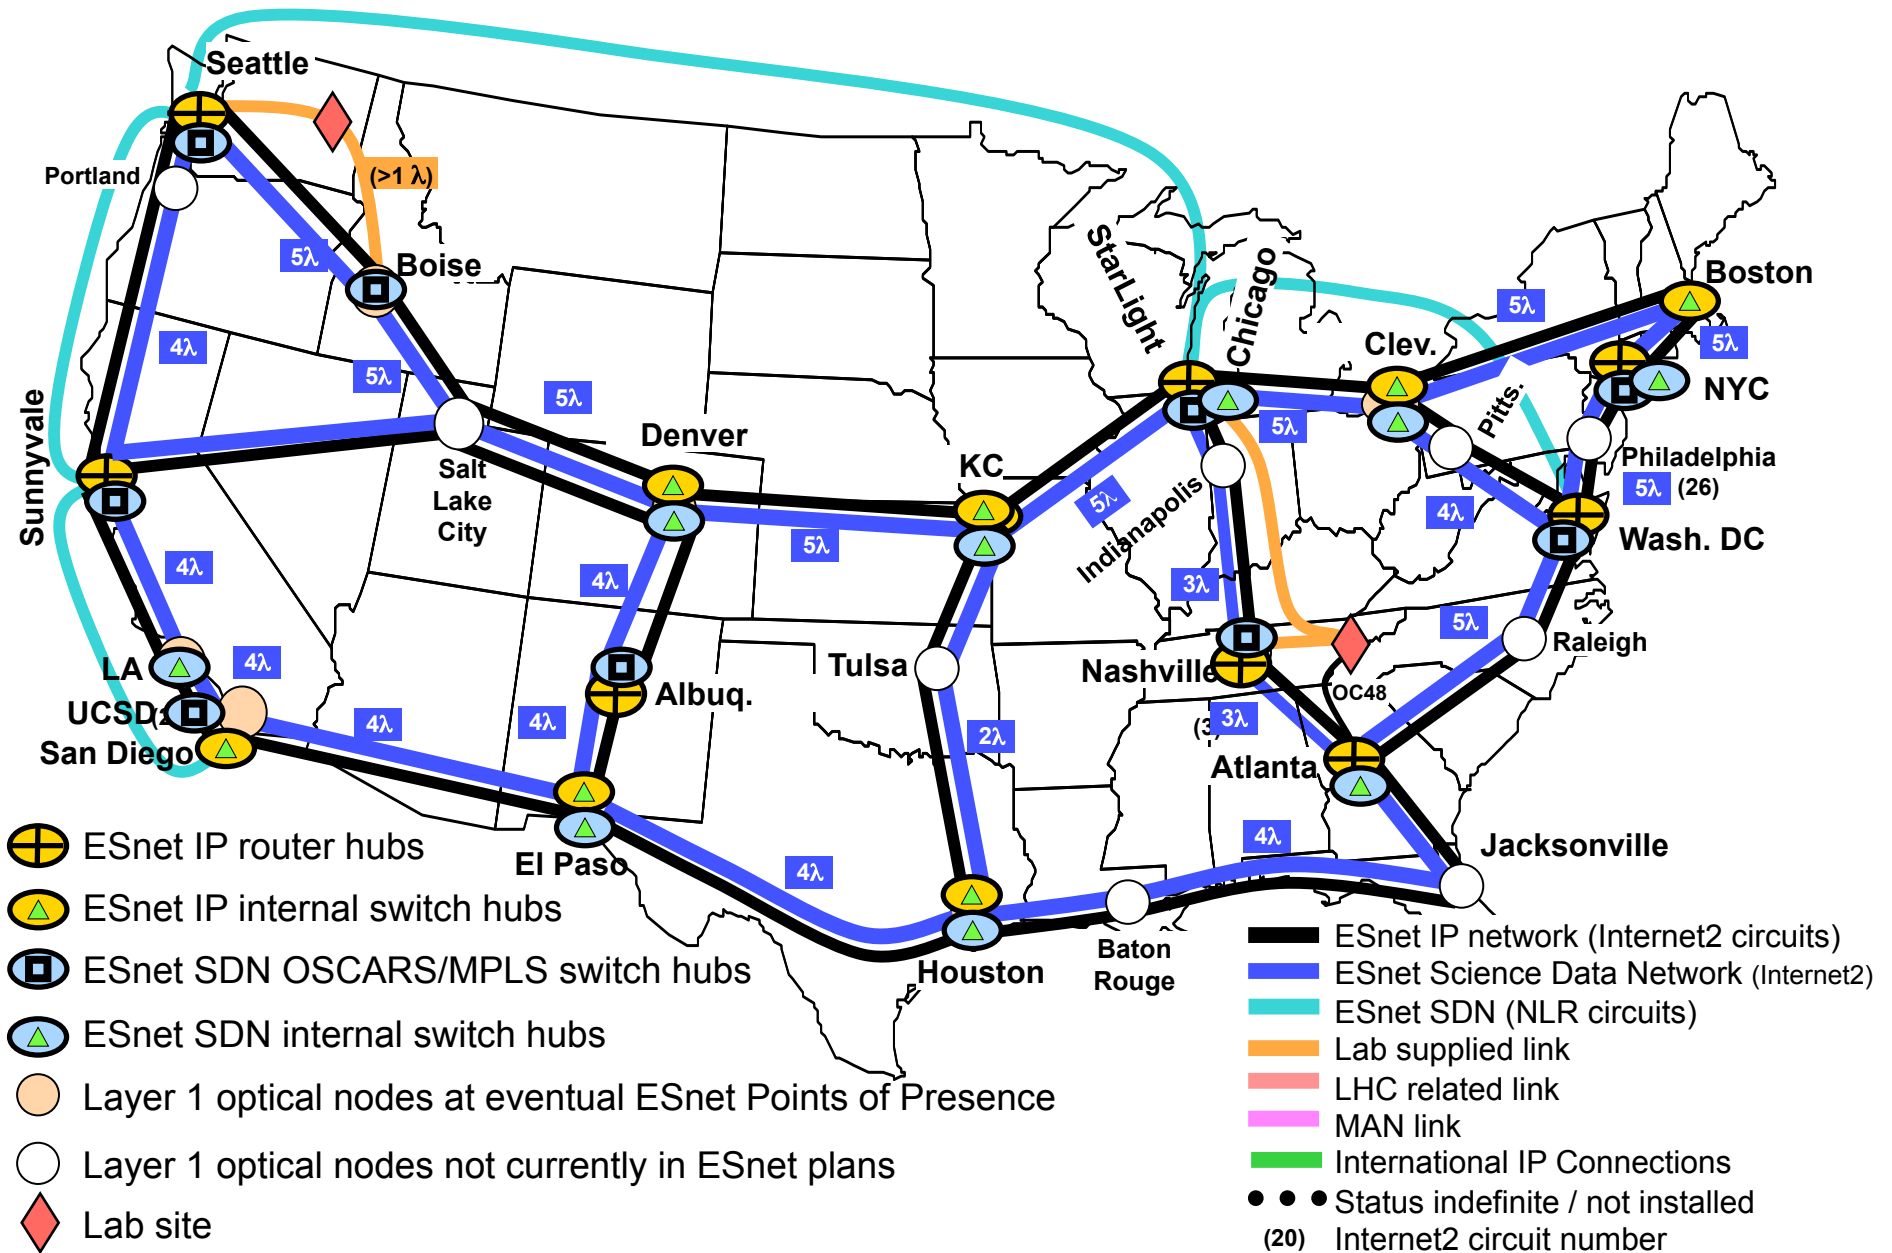
Layer 3 VC Service:
Packets matching reservation profile IP flow-spec are filtered out (i.e. policy based routing), "policed" to reserved bandwidth, and injected into an LSP.

Layer 2 VC Service:
Packets matching reservation profile VLAN ID are filtered out (i.e. L2VPN), "policed" to reserved bandwidth, and injected into an LSP.

LSP between ESnet border routers is determined using topology information from OSPF-TE. Path of LSP is explicitly directed to take SDN network where possible. On the SDN Ethernet switches all traffic is MPLS switched (layer 2.5).



ESnet4 IP + SDN, 2011 Configuration (Est.)



OSCARS Status Update

- **ESnet Centric Deployment**
 - Prototype layer 3 (IP) guaranteed bandwidth virtual circuit service deployed in ESnet (1Q05)
 - Prototype layer 2 (Ethernet VLAN) virtual circuit service deployed in ESnet (3Q07)
 - Support soft reservations (2Q08)
 - Automatic graph generation of VCs (2Q08)
 - Support site administrator role (2Q08)
- **Inter-Domain Collaborative Efforts**
 - **Terapaths**
 - Inter-domain interoperability for layer 3 virtual circuits demonstrated (3Q06)
 - Inter-domain interoperability for layer 2 virtual circuits demonstrated at SC07 (4Q07)
 - **LambdaStation**
 - Inter-domain interoperability for layer 2 virtual circuits demonstrated at SC07 (4Q07)
 - **I2 DCN/DRAGON**
 - Inter-domain exchange of control messages demonstrated (1Q07)
 - Integration of OSCARS and DRAGON has been successful (1Q07)
 - **GEANT2 AutoBAHN**
 - Inter-domain reservation demonstrated at SC07 (4Q07)
 - **DICE**
 - First draft of topology exchange schema has been formalized (in collaboration with NMWG) (2Q07), interoperability test demonstrated 3Q07
 - Initial implementation of reservation and signaling messages demonstrated at SC07 (4Q07)
 - **Nortel**
 - Topology exchange demonstrated successfully 3Q07
 - Inter-domain interoperability for layer 2 virtual circuits demonstrated at SC07 (4Q07)
 - **UVA**
 - Demonstrated token based authorization concept with OSCARS at SC07 (4Q07)
 - **OGF NML-WG**
 - Actively working to combine work from NMWG and NDNL
 - Documents and UML diagram for base concepts have been drafted (2Q08)
 - **GLIF GNI-API WG**
 - In process of designing common API and reference middleware implementation

OSCARS Operational Issues

- Site Managed Load-Balancing

E.g.

- Site wants 5 L2 VCs for 5 distinct site-to-site connections on a 10GE path
- Some of the VCs share common link elements in the path
- Each VC is capable of bursting up to 10Gb/s, but concurrently use of all 5 VCs are constrained by shared 10GE link element
- To provision guaranteed bandwidth in OSCARS, cumulative bandwidth request for all 5 L2 VCs cannot exceed 10Gb/s (over-provisioning is not allowed)

- Redundant Backup VC

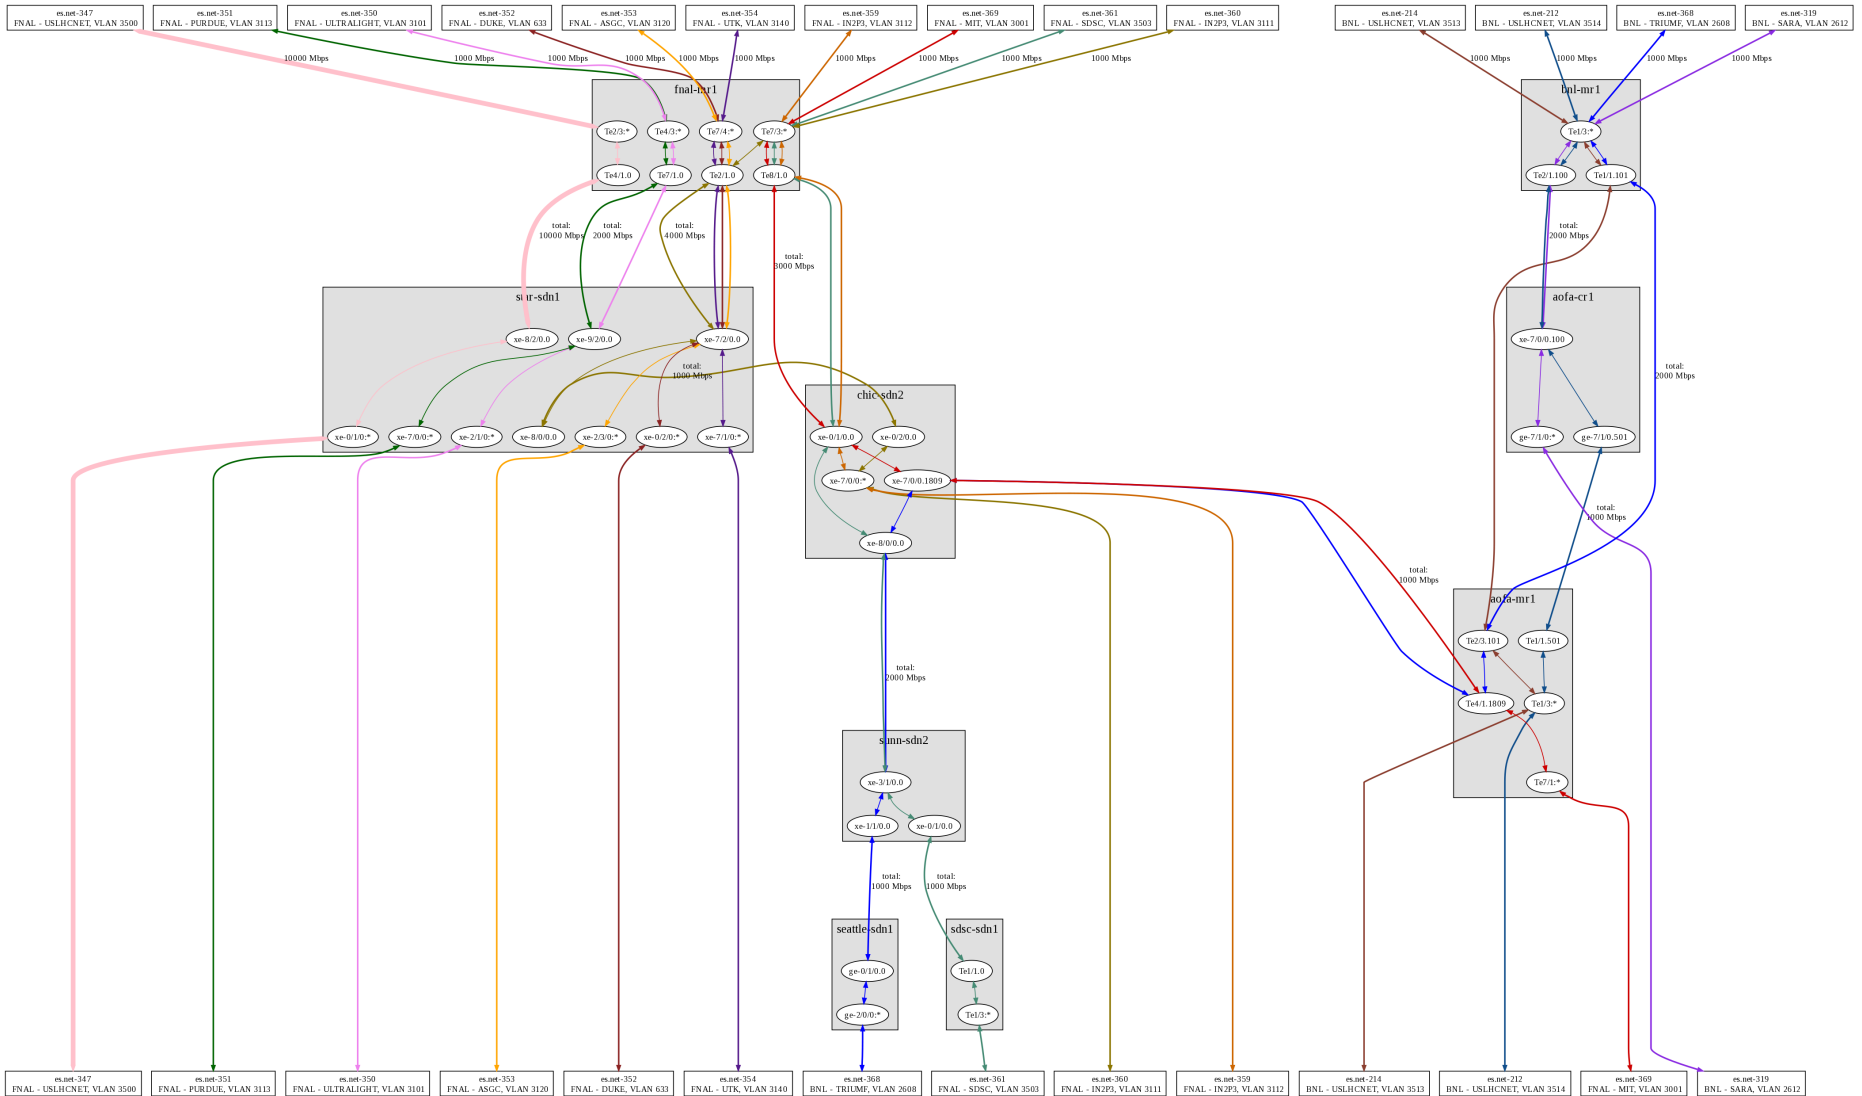
E.g.

- Site wants 2 L2 VCs on 10GE path, one for primary, and one for backup
- Both VCs share common link elements in path
- Only one of the the VCs will be in use at any time
- Both VCs cannot be requested at 10Gb/s because they are viewed as distinct reservations by OSCARS

OSCARS Operational Changes

- Previous Configuration
 - Guaranteed Bandwidth VCs
 - Over-provisioning is prevented at reservation request time
 - Over-subscription is prevented by policing (hard drop) at time of use
 - VC is configured to transit ESnet as Expedited Forwarding Class
- Current Configuration
 - Guaranteed Bandwidth VC with Over-Subscription
 - Over-provisioning is prevented at reservation request time
 - Over-subscription is allowed at transfer time
 - Traffic below policed rate will transit ESnet as Expedited Forwarding Class
 - Traffic above policed rate is not dropped, but remarked as Scavenger Class
 - Considerations
 - Implementation of above enhancements are technology specific
 - End-to-end inter-domain dynamic VCs may not support over-subscription
 - Multi-lab coordination may be required to effectively utilize bandwidth available in Scavenger Class

OSCARS Production Circuits (as of 20080714)



OSCARS: Guaranteed Bandwidth Service

- Funded by the DOE Office of Science
- Info URL: <http://www.es.net/oscars>
- Services URL: <https://oscars.es.net/OSCARS/>
- Contact:
 - Chin Guok (chin@es.net)
 - David Robertson (dwrobertson@lbl.gov)
 - Evangelios Chaniotakis (haniotak@es.net)