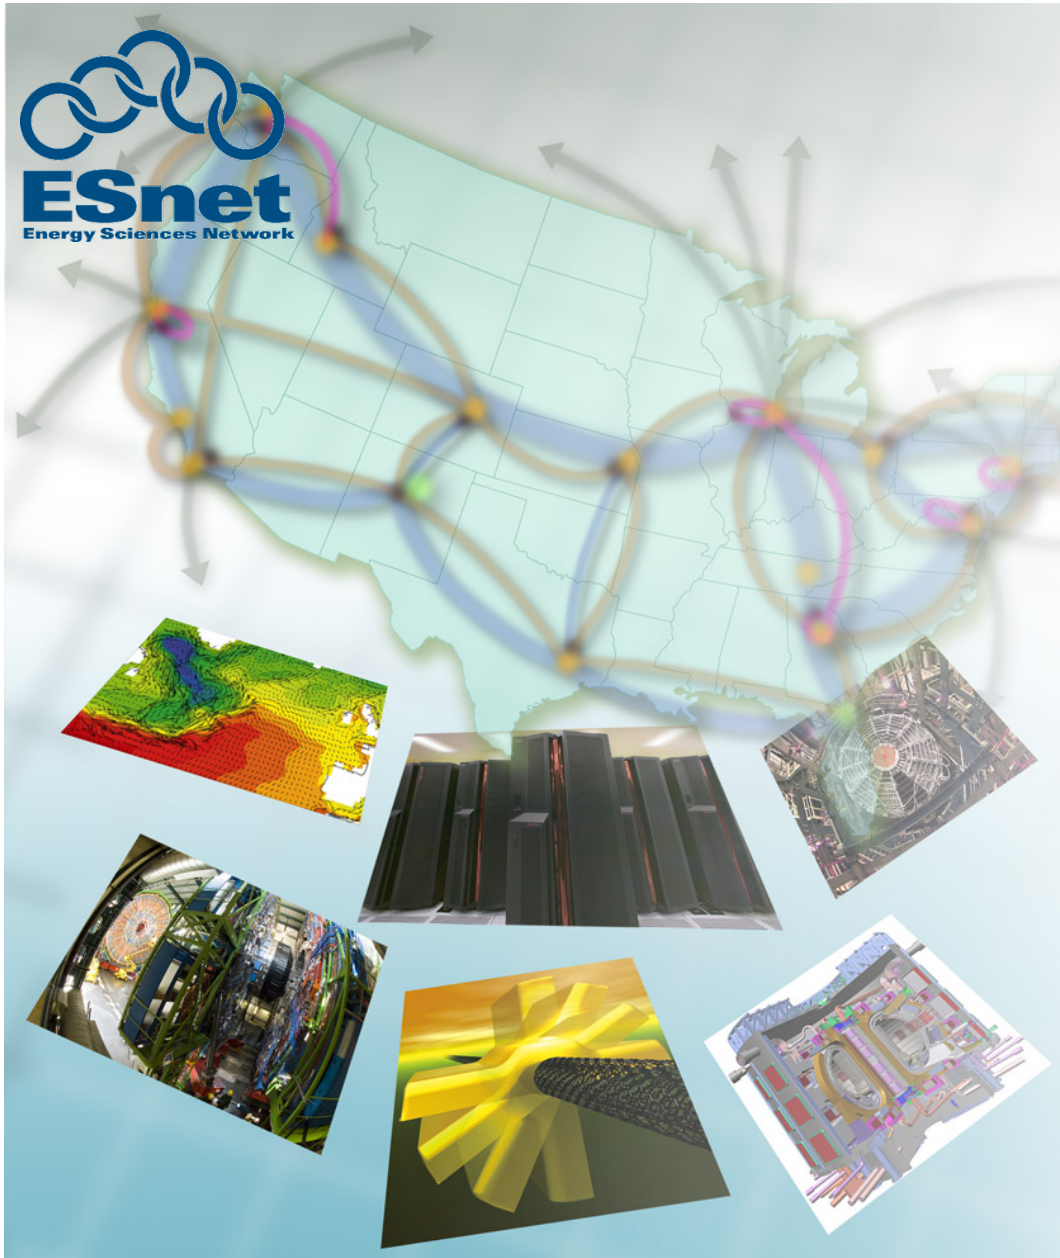

OSCARS Roadmap

Chin Guok (chin@es.net)

Energy Sciences Network
Lawrence Berkeley National Laboratory

Feb 6, 2009

Networking for the Future of Science

OSCARS Status

- Development of a community approach that will support end-to-end virtual circuits in the R&E environment is coordinated by the DICE (Dante, Internet2, Caltech, ESnet) working group (which involves many more organizations than the original group)
  - Each organization potentially has their own InterDomain Controller approach (though the ESnet/Internet2 OSCARS code base is used by several organizations (flagged OSCARS/DCN)
  - The DICE group has developed a standardized InterDomain Control Protocol (IDCP) for specifying the set up of segments of end-to-end VCs
    - While there are several very different InterDomain Controller implementations, they all speak IDCP and support compatible data plane connections
  - The following organizations have *implemented/deployed systems which are compatible with the DICE IDCP*:
    - Internet2 Dynamic Circuit Network (OSCARS/DCN)
    - ESNet Science Data Network (OSCARS/SDN)
    - GÉANT2 AutoBahn System
    - Nortel (via a wrapper on top of their commercial DRAC System)
    - Surfnet (via use of above Nortel solution)
    - LHCNet (OSCARS/DCN)
    - Nysernet (New York RON) (OSCARS/DCN)
    - LEARN (Texas RON) (OSCARS/DCN)
    - LONI (OSCARS/DCN)
    - Northrop Grumman (OSCARS/DCN)
    - University of Amsterdam (OSCARS/DCN)
    - DRAGON (U. Maryland/MAX) Network
  - The following "higher level service applications" have adapted their existing systems to communicate via the user request side of the IDCP:
    - LambdaStation (FermiLab)
    - TeraPaths (Brookhaven)
    - Phoebus (UMd)

OSCARS Production VCs (as of 01/27/2009)

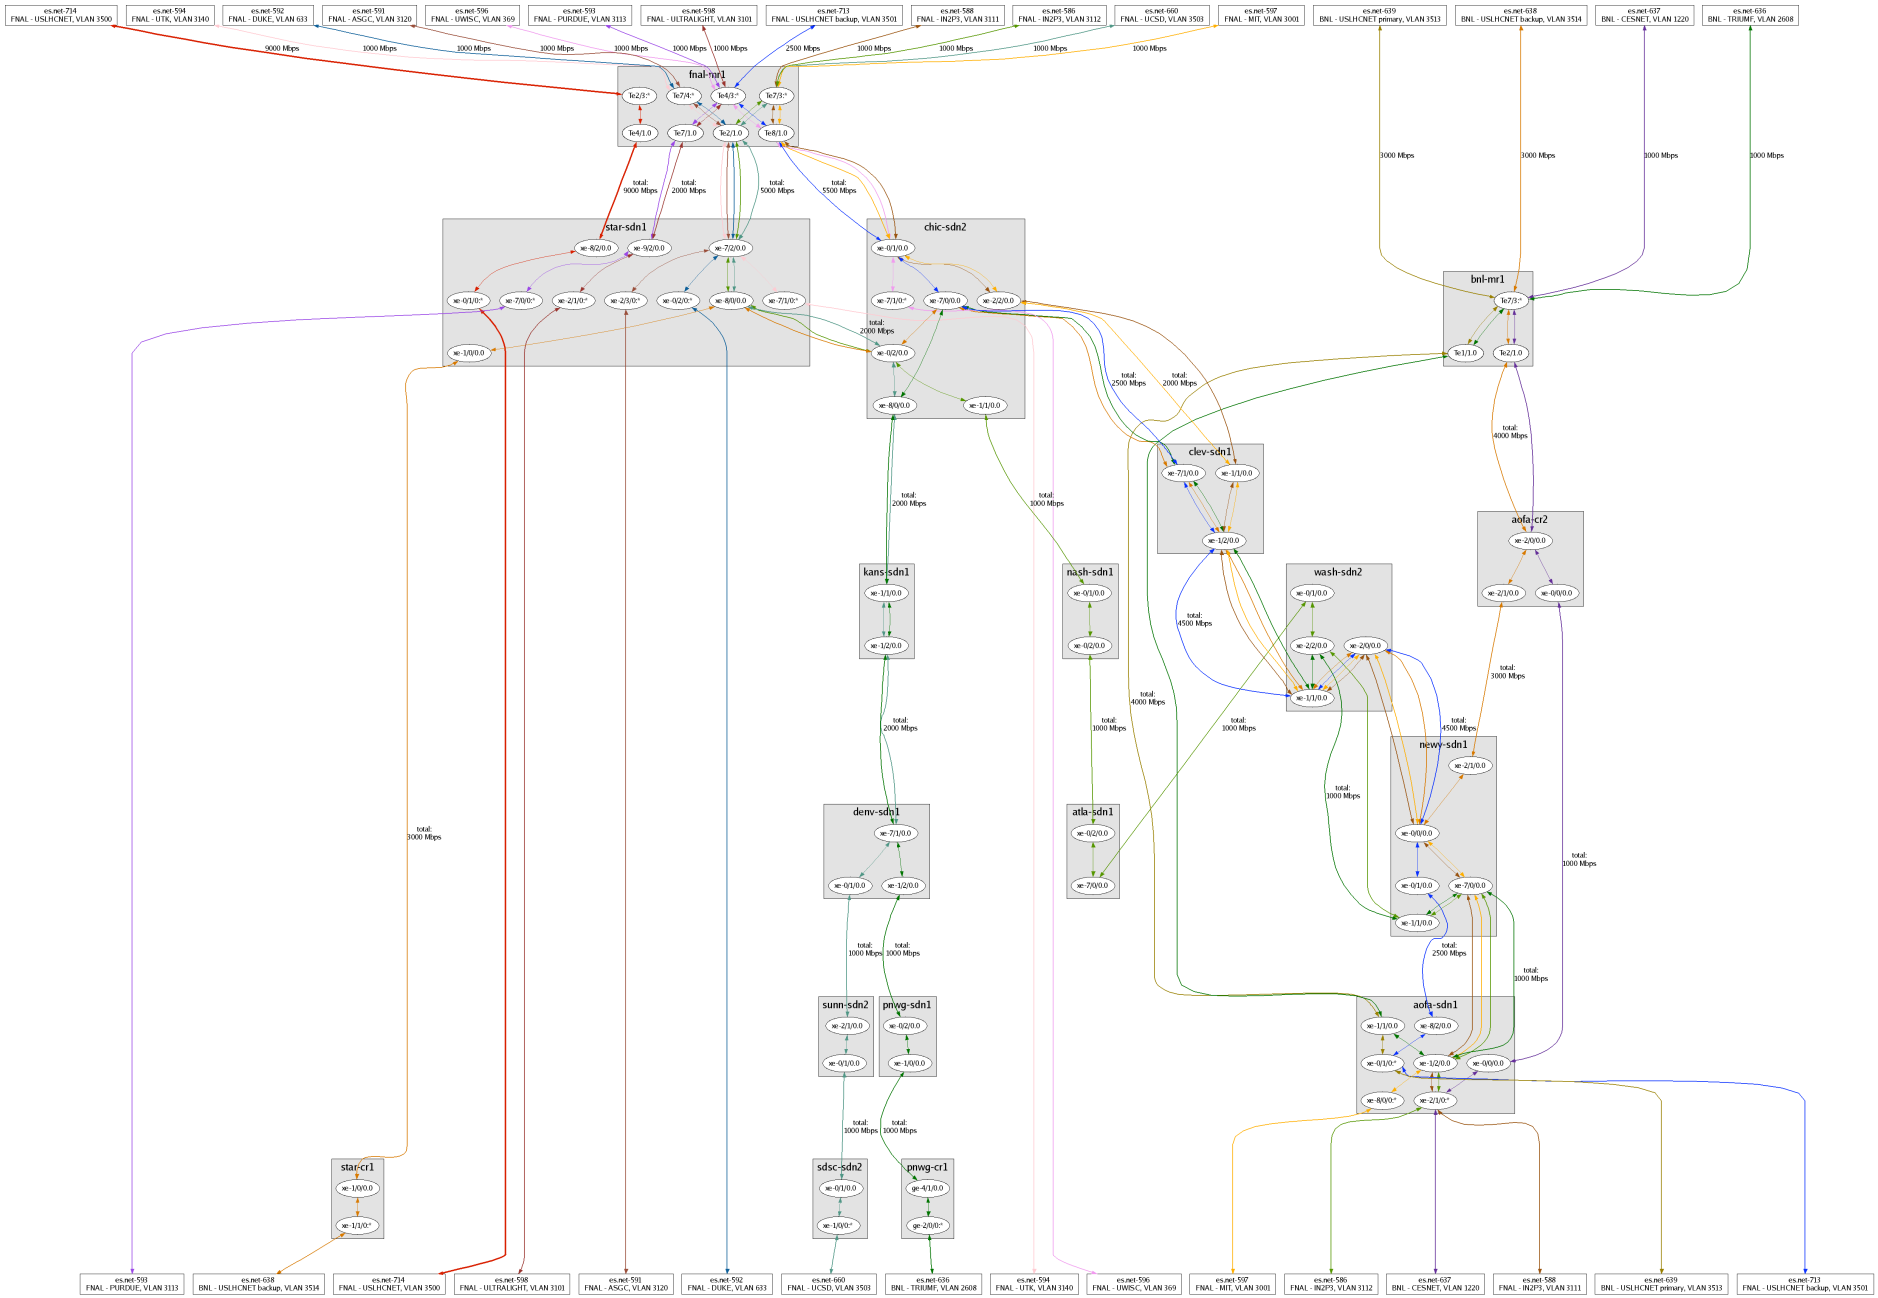

OSCARS Evolution

- Phase 1
  - Proof of concept
  - Intra-domain virtual circuit (VC) services
- Phase 2
  - Inter-domain interoperability
  - Pre-production ESnet VC services
- Phase 3 (Current)
  - Productionalizing OSCARS (IDC)
- Phase 4
  - Extending service offerings
  - Framework for research
  - *Explore "on-ramp" technologies*

Network 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).

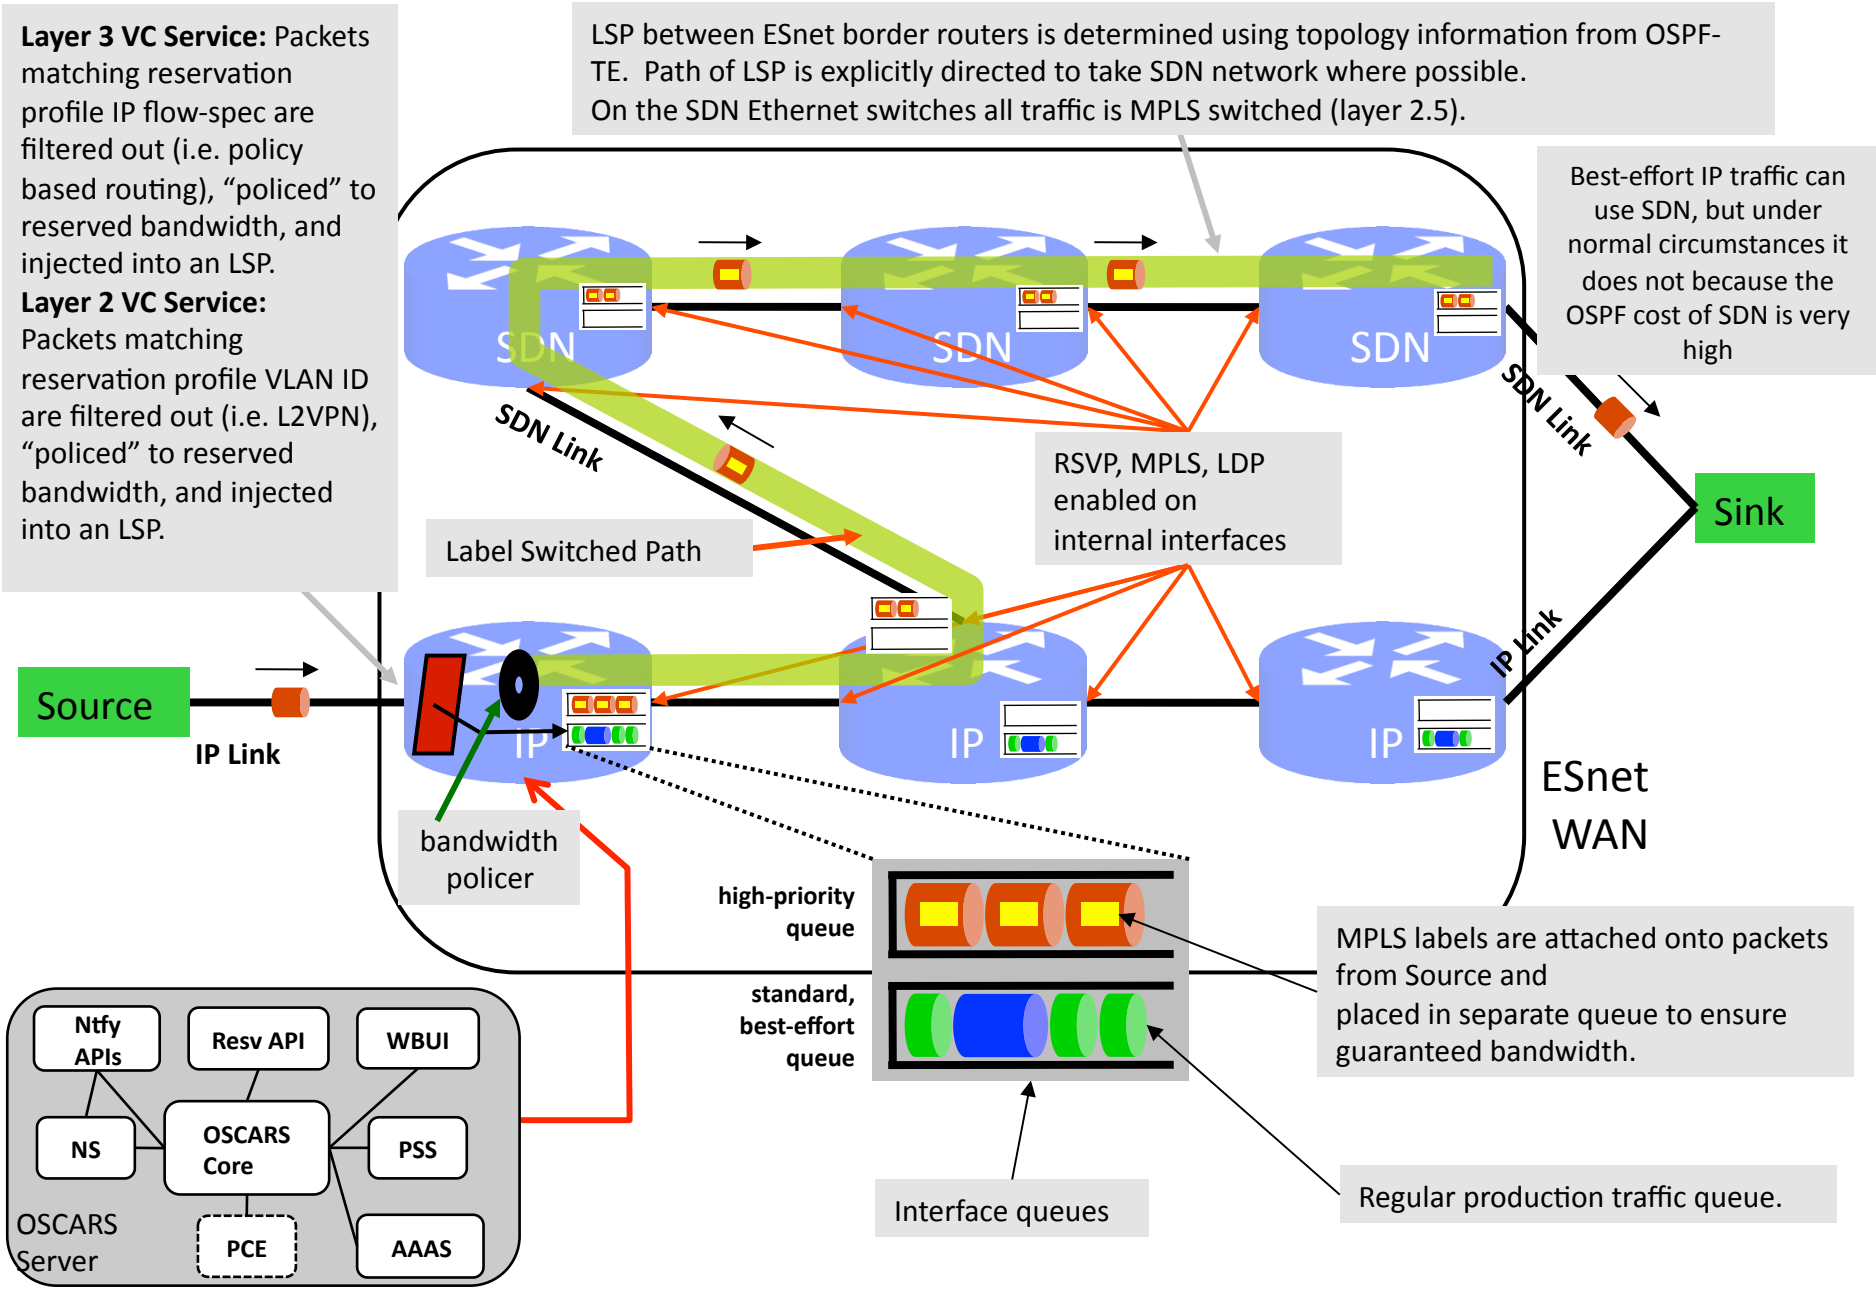

OSCARS Initial Implementation

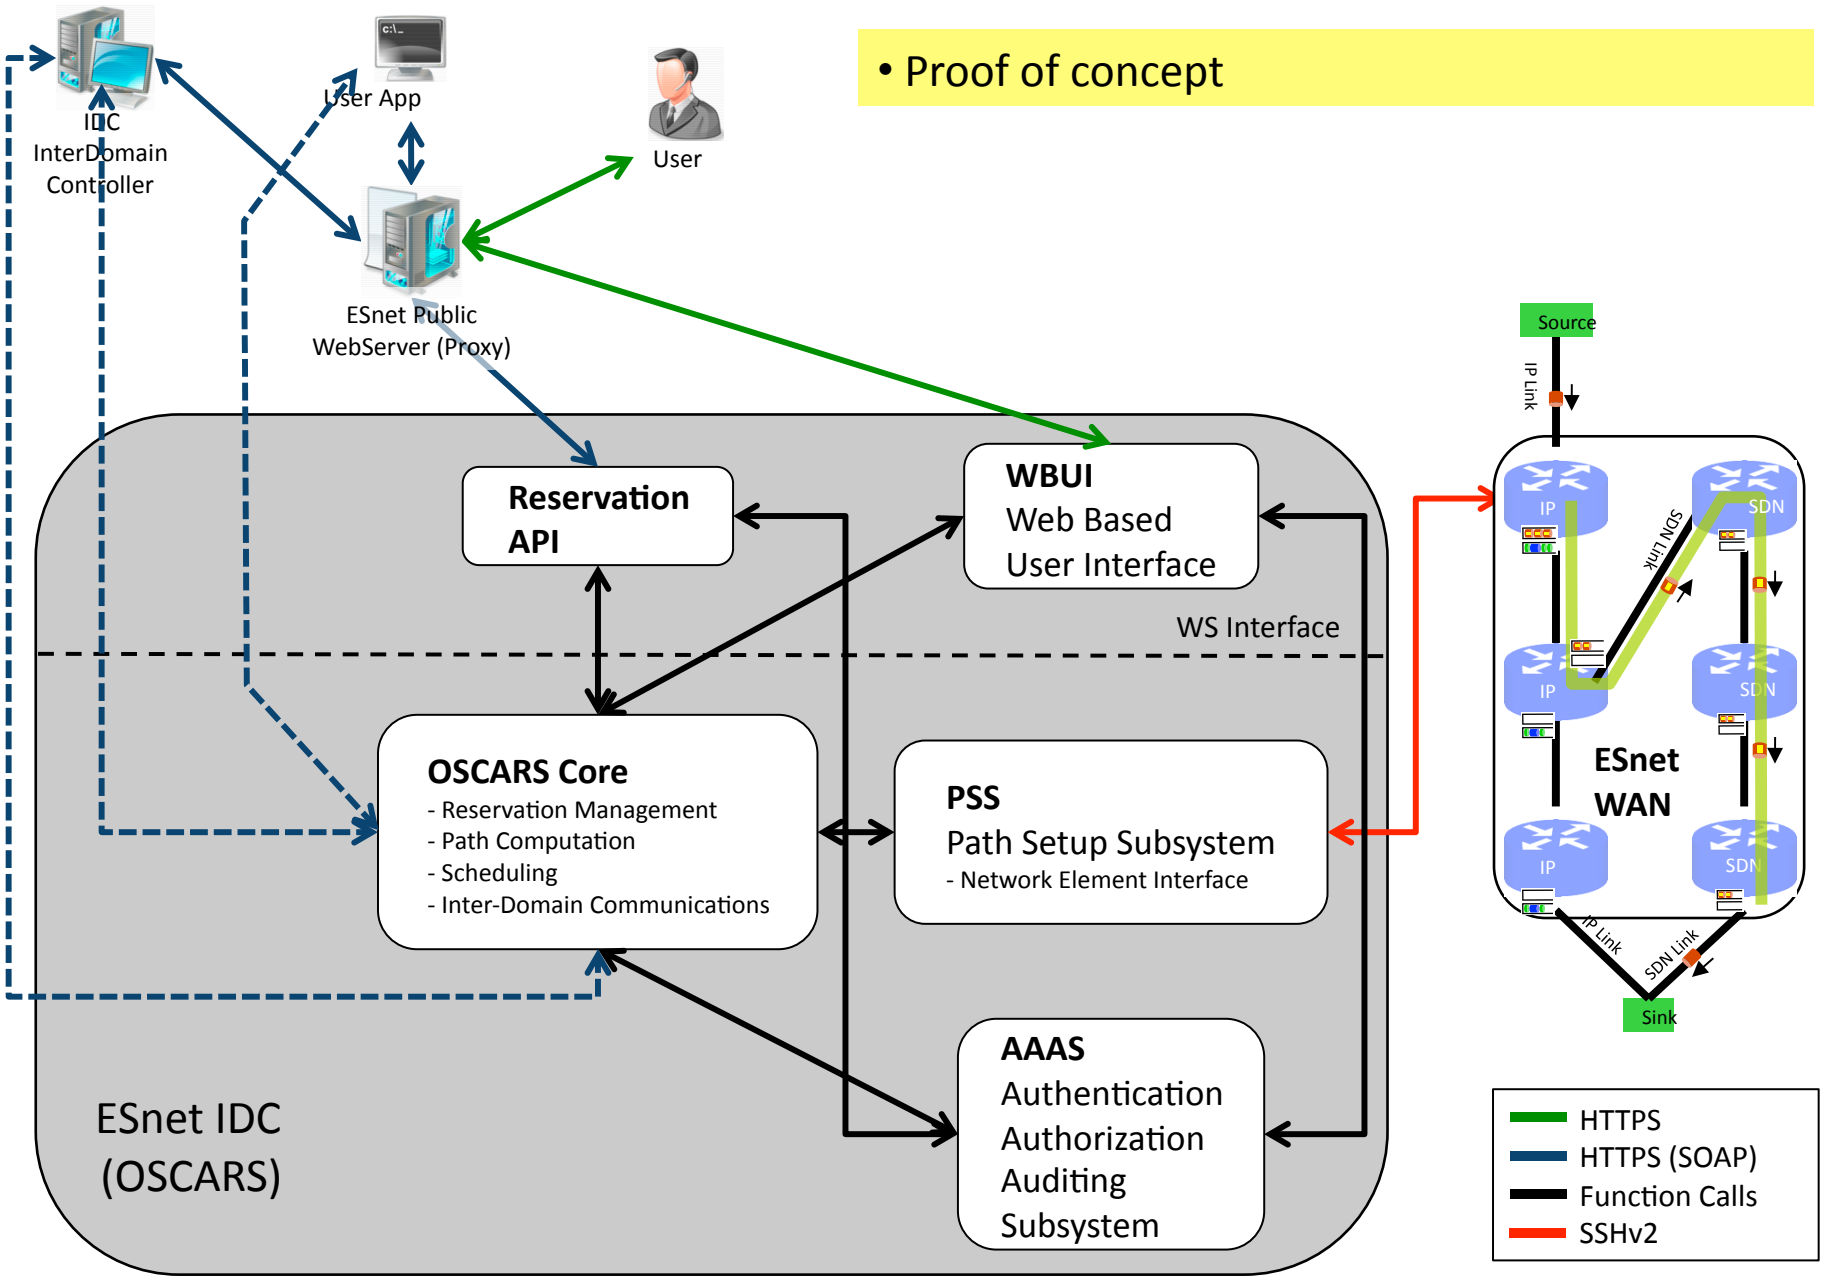

OSCARS Current (ver 0.5) Implementation (2Q09)

- Well defined inter-module interfaces
- Exchange of static topology information
- PCE integrated into OSCARS Core

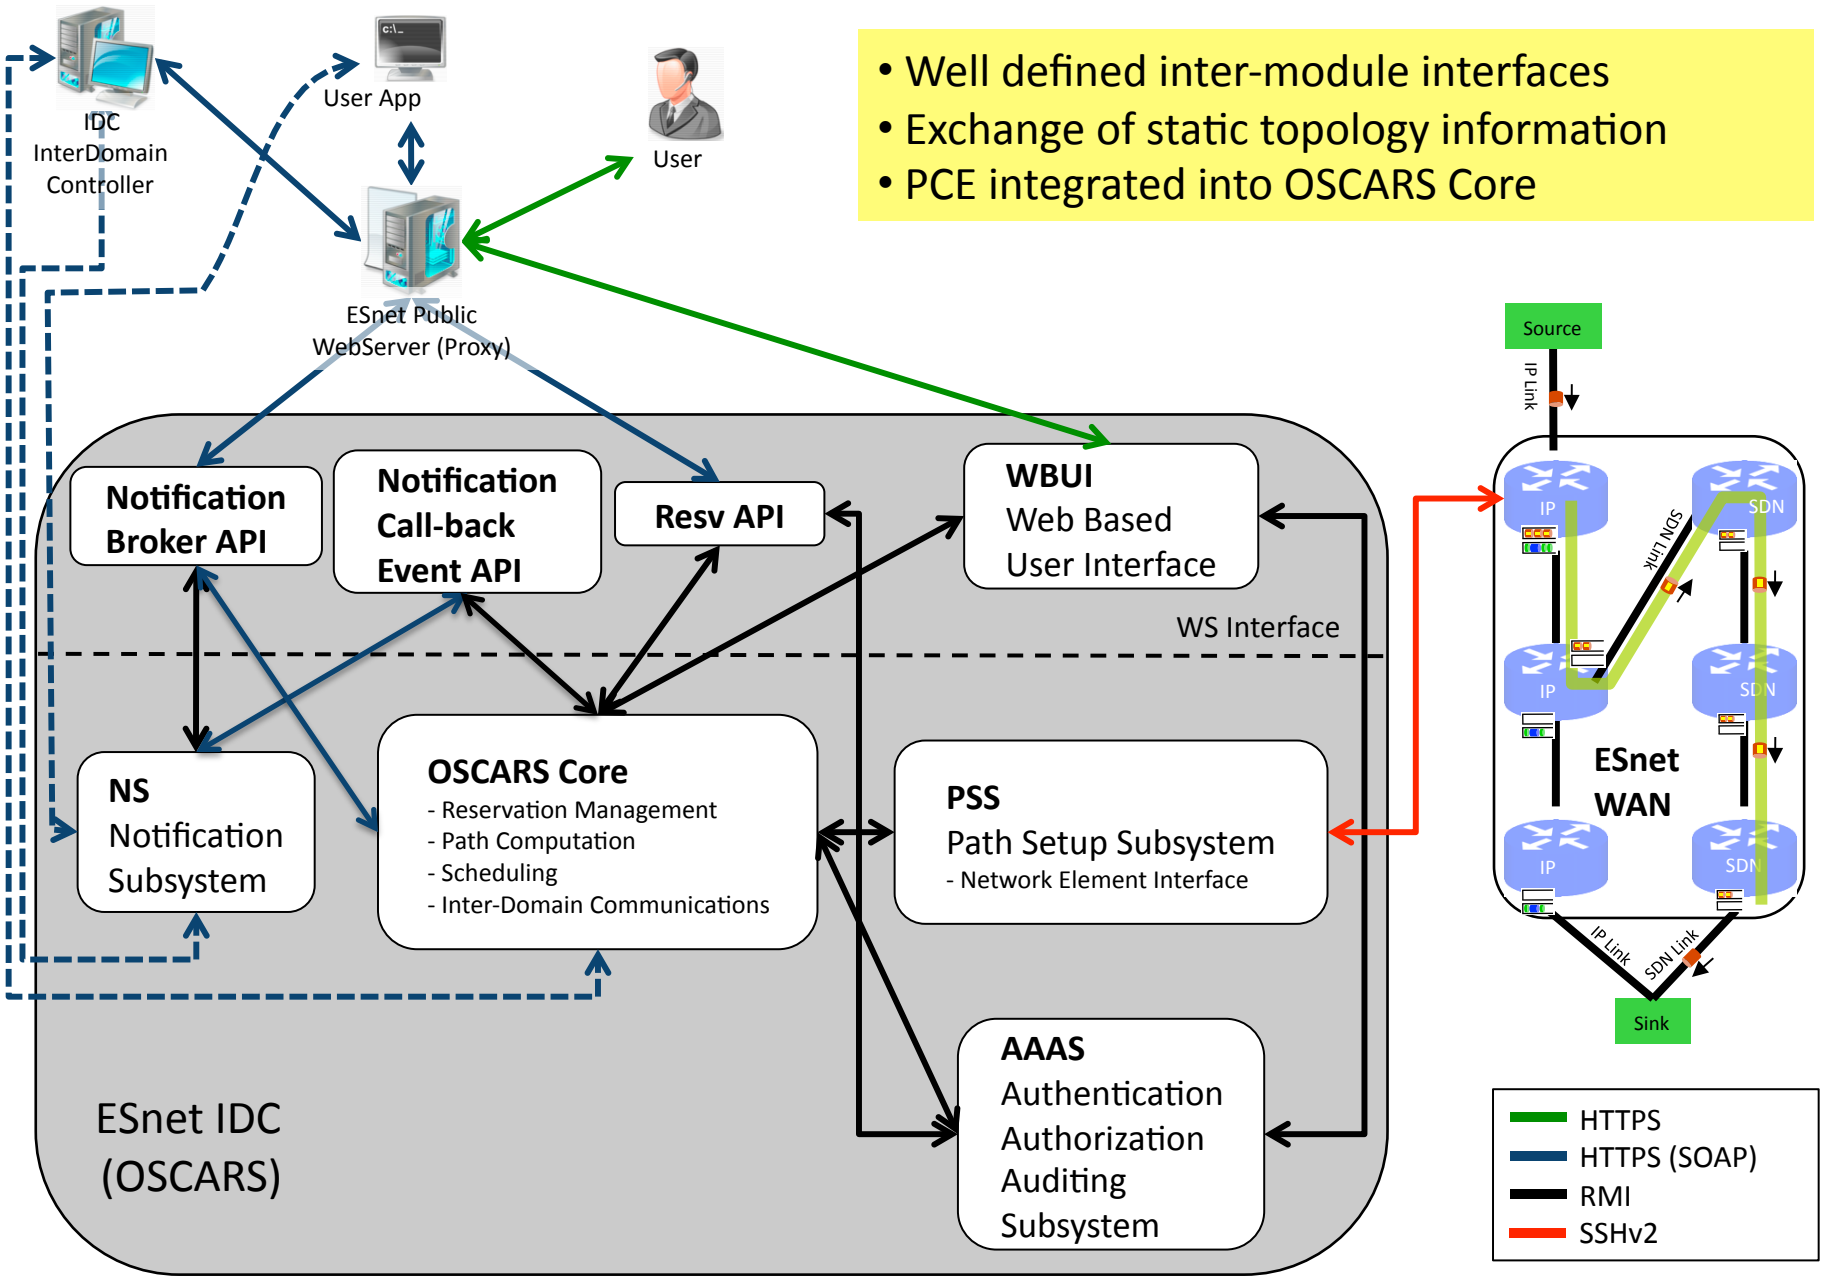

OSCARS Future Implementation (4Q09)

- Exchange of dynamic topology information
  - includes time dimension
- PCE separated from OSCARS Core
  - PCEs can be daisy changed
  - allows PCE to be pluggable
  - facilitates a research framework for collaboration

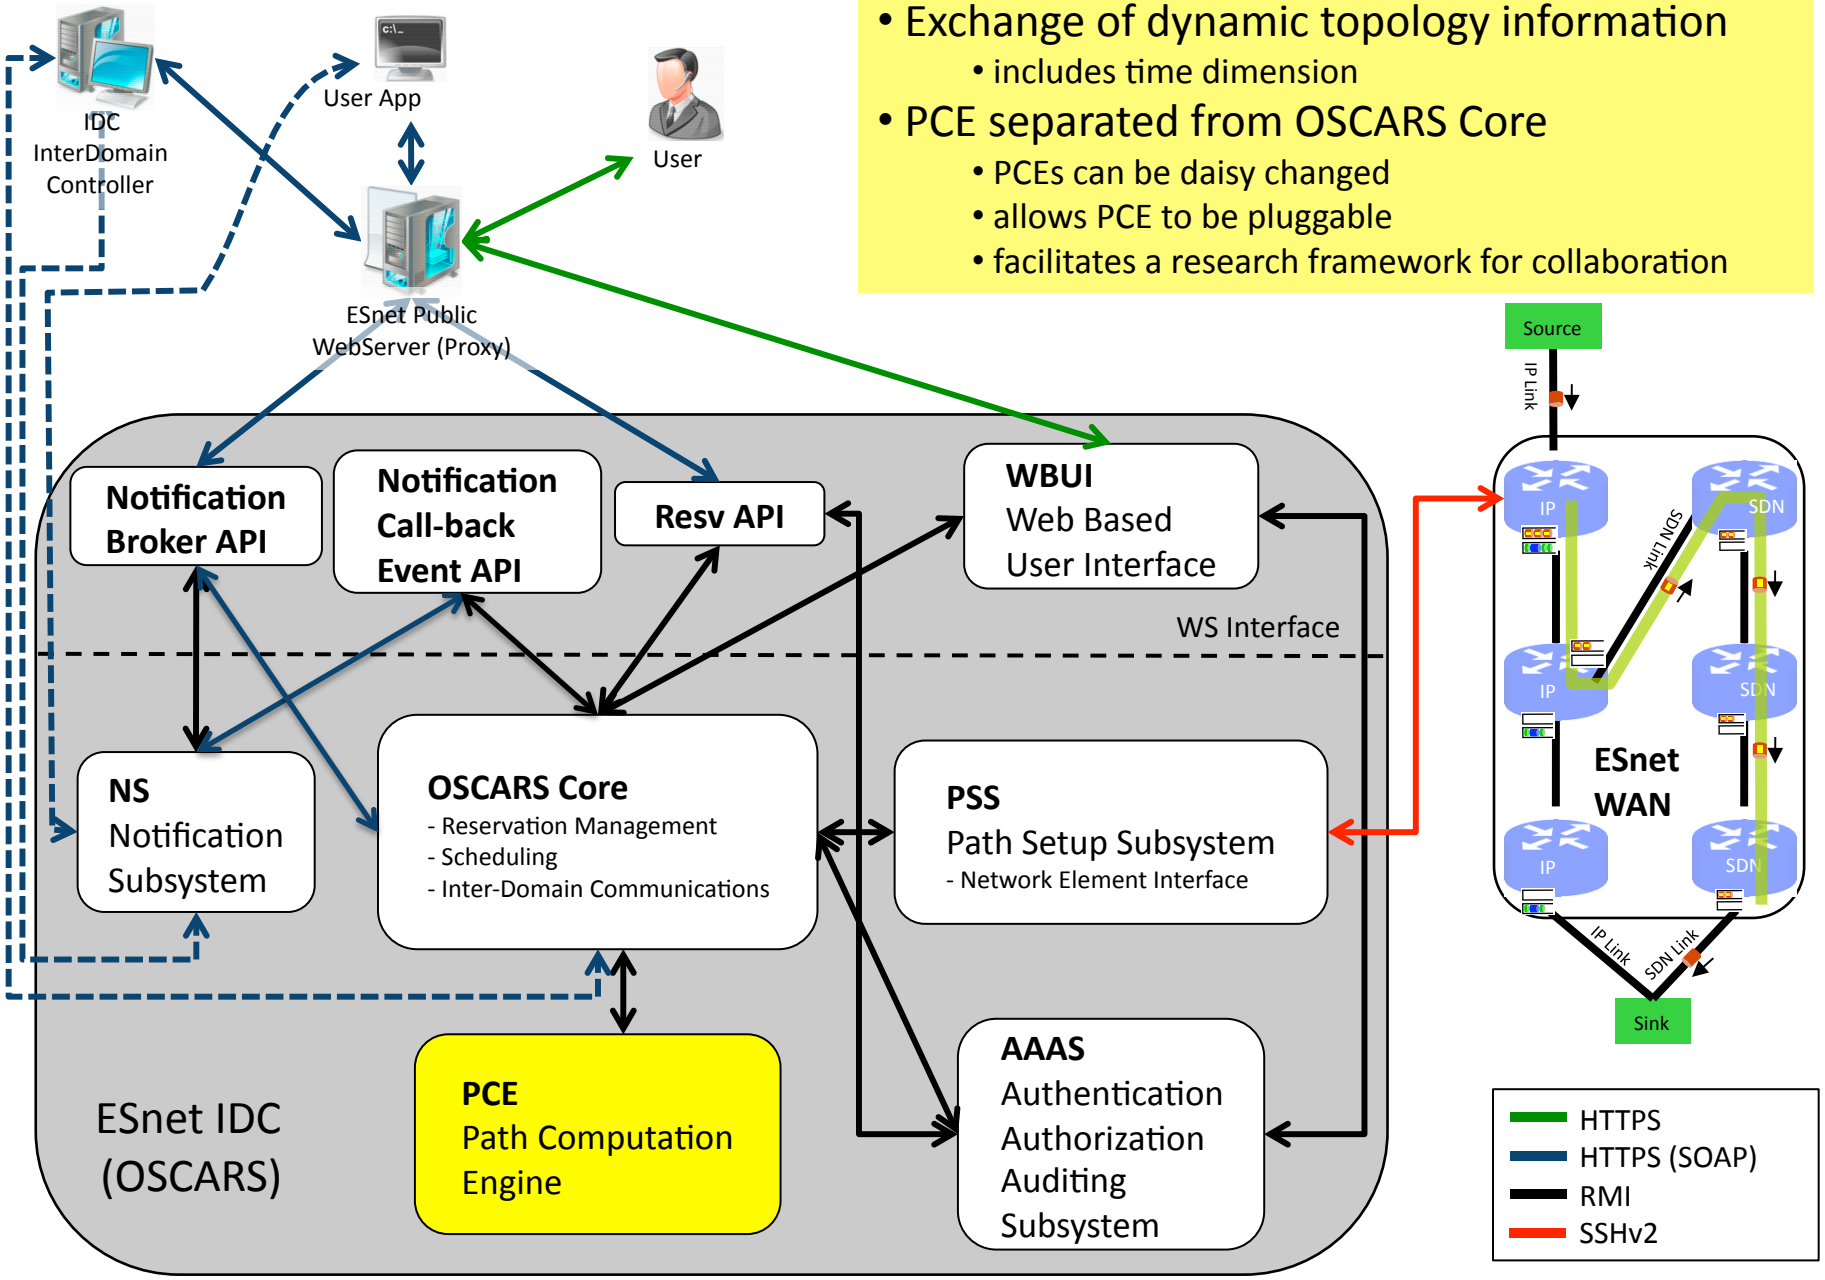

SDN "On Ramp" Options

- There are a number of ways to move flows from the ESnet IP cloud to the SDN
- Depends on the following factors
  - Control: ESnet admin or site admin or user
  - Advanced reservations or on-demand
  - Manual vs. automatic
  - End-to-end vs. border to border
  - Level of expertise needed
  - Layer 1 vs layer 2 vs layer 3 circuit

Example: ESnet traffic engineering to avoid hot-spots

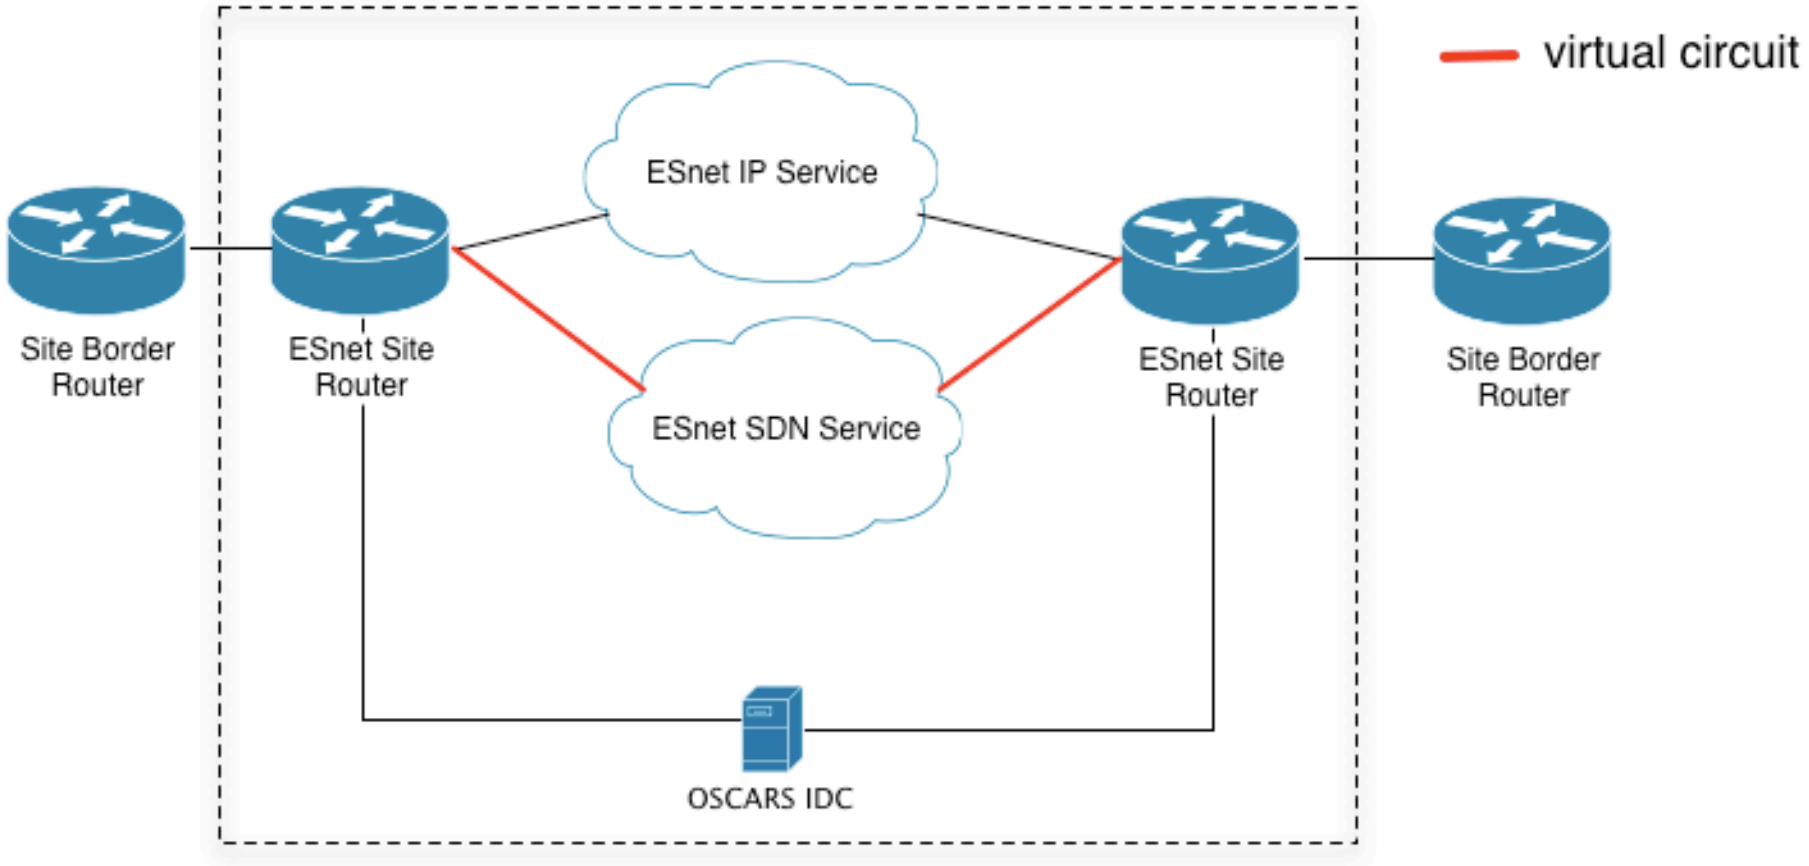

- Off-load congested IP circuits and traffic engineer around hot-spots
- Transparent to site networking
- Detailed monitoring data of ESnet traffic is needed to determine optimal solution

Example: LHC Tier 1 – Tier 2 data movement

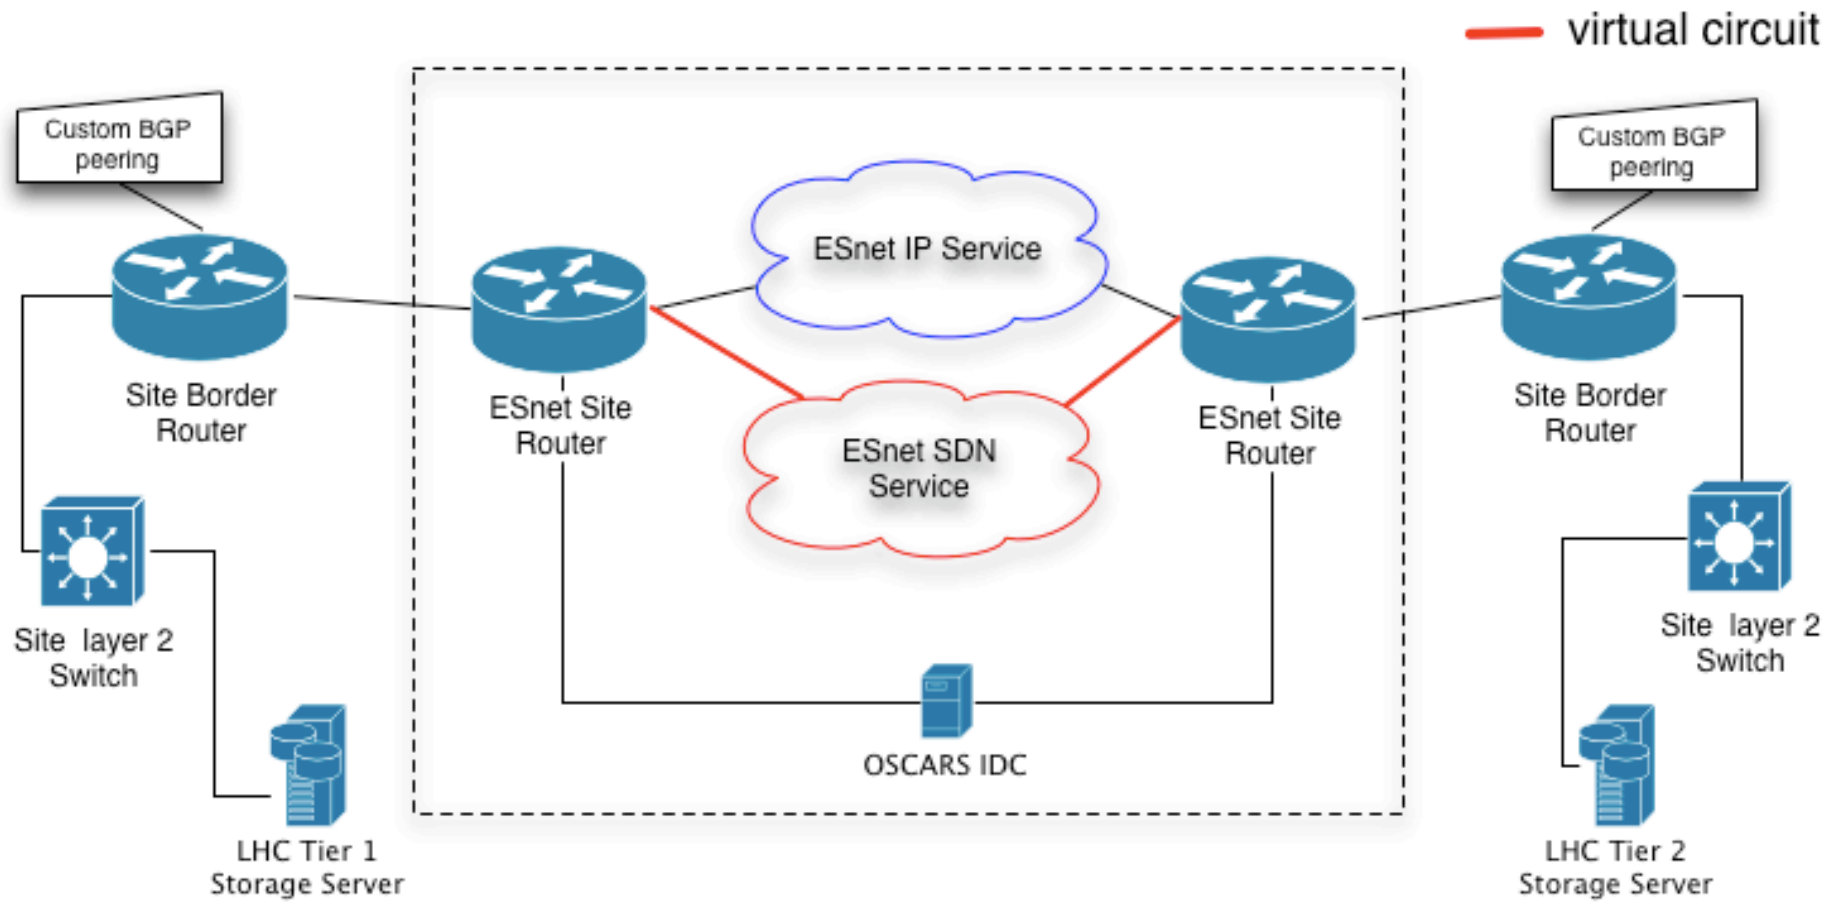

- Site configures BGP between CE routers to use SDN VC if available
- User initiates request to OSCARS for SDN VC

Example: Fusion Simulation at NERSC writing results to ORNL over GPFS

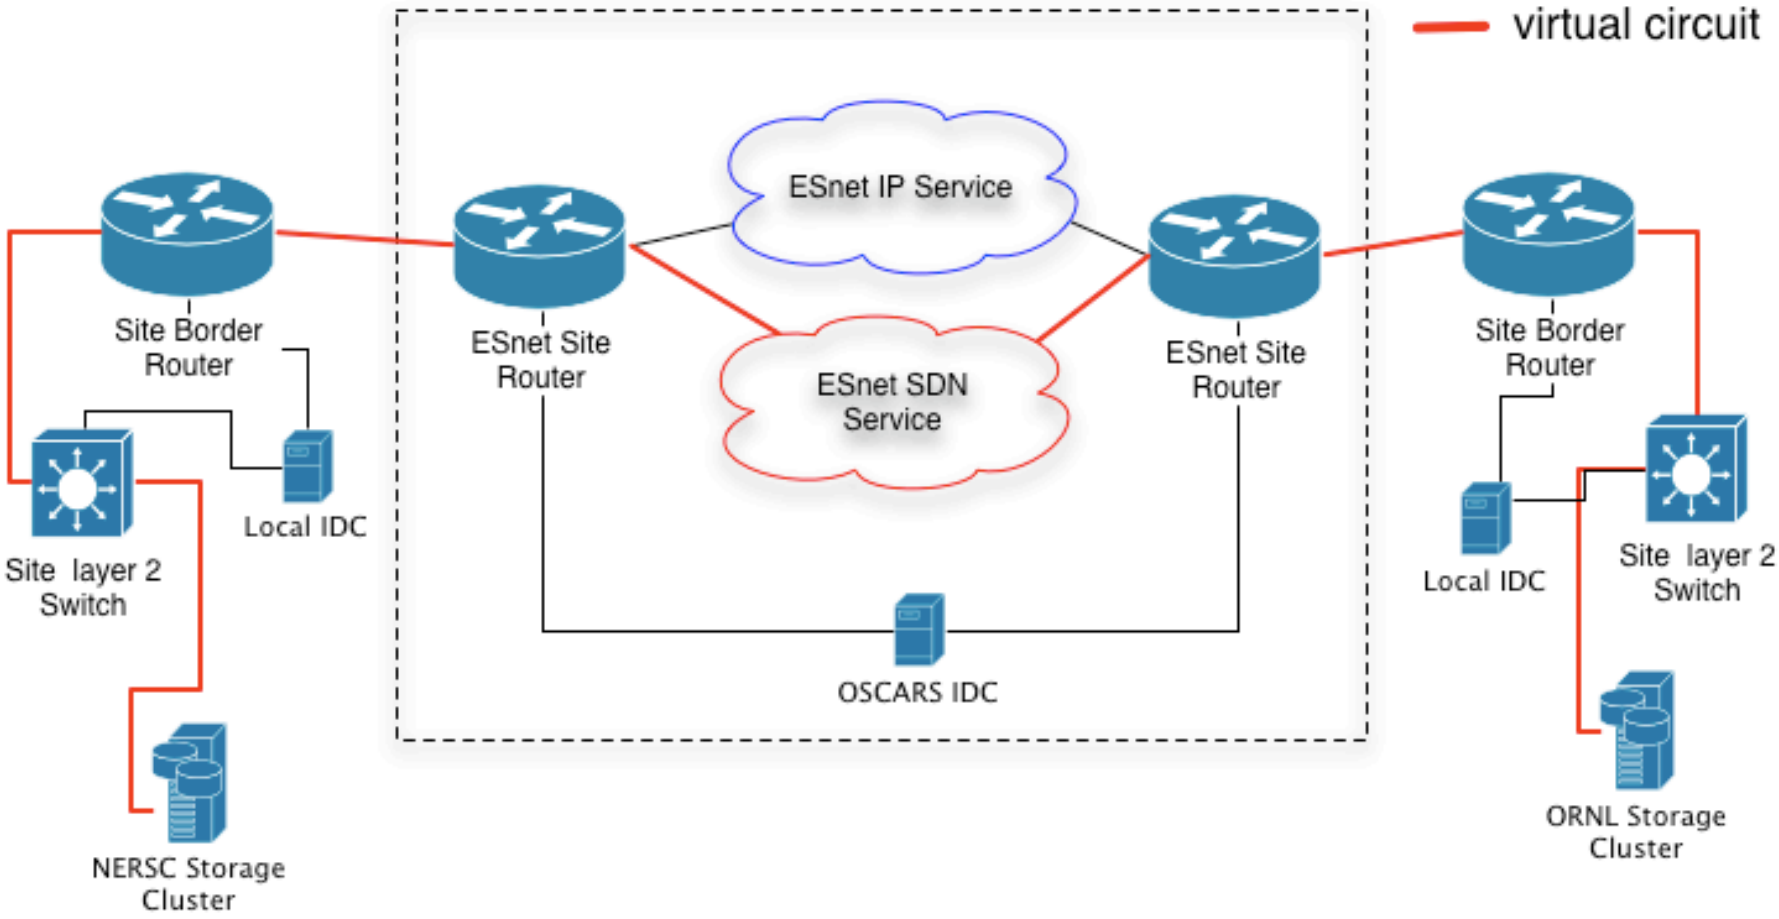

- Site admin/user requests VC from any IDC
- Local domain controller (e.g. OSCARS/LambdaStation/Terapaths) performs VC setup
- Typically scheduled in advance

Example: Climate data movement from ANL to LLNL

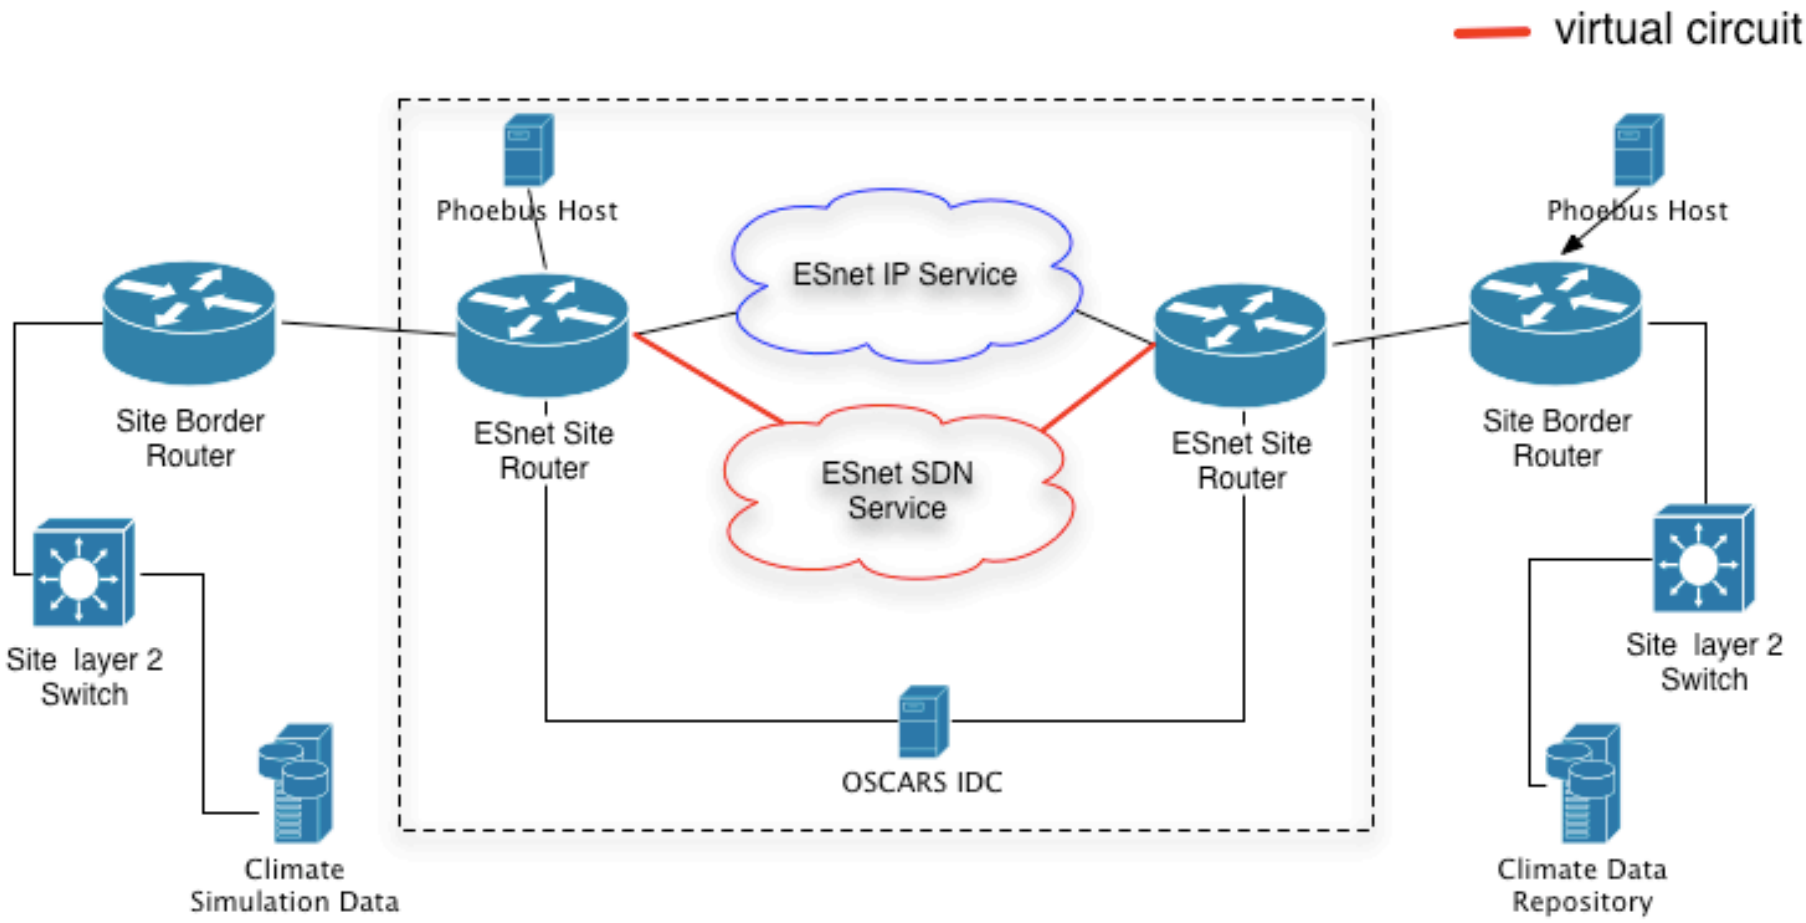

- User points to closest Phoebus host
- Use Phoebus to forward data to SDN
- Phoebus host requests OSCARS for SDN VC on behalf of the user
- Phoebus host could be part of ESnet or site infrastructure

Example: APS user sending data to home institute

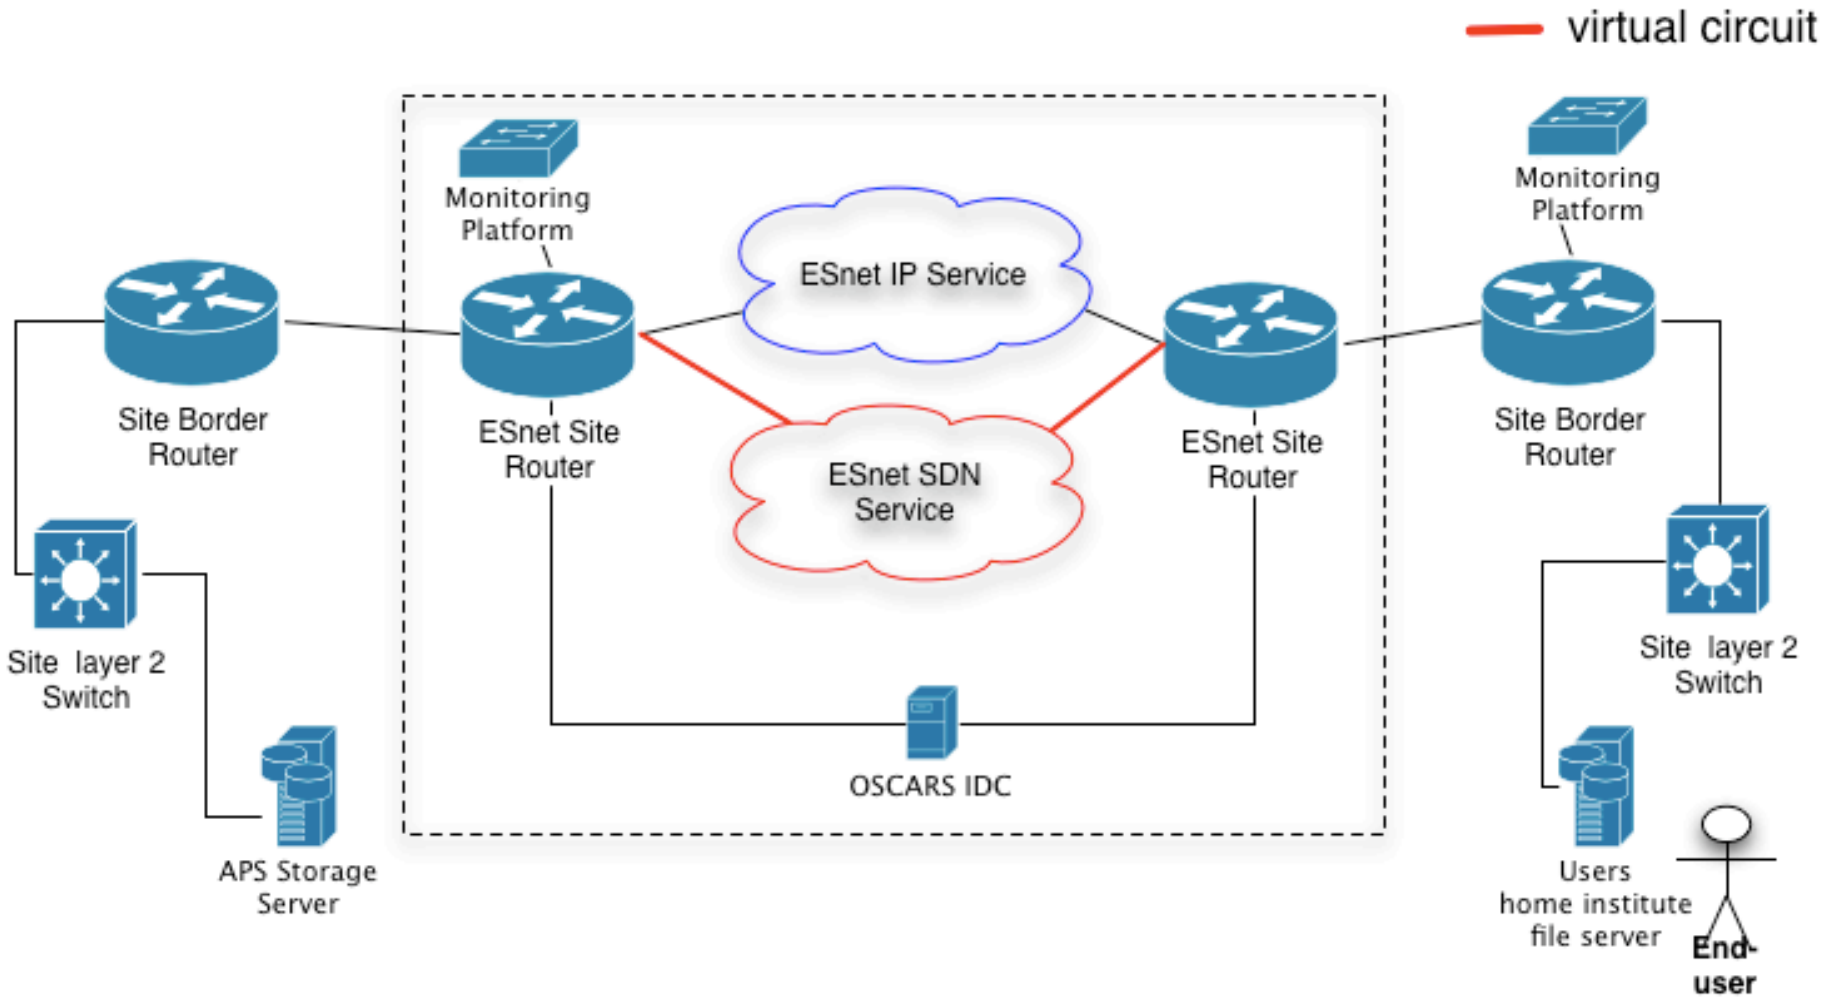

- Router forwards packets to monitoring device (e.g. Juniper MS PIC, Bro host, etc)
- Monitoring device identifies candidate flows for SDN, requests SDN VC
- Monitoring devices could be part of ESnet or part of site infrastructure