


Network Stitching

GEC5 Control Framework Working Group

Chris Tracy
Jarda Flidr
Peter O'Neil
Cluster B Participant
July 22nd, 2009



Outline

- Goals of Network Stitching
- Our view on Deterministic Services
- Challenges of Network Stitching
- End-to-End Slices Across Network Aggregates
 - Ontology adopted by DCN Software Suite (DRAGON/OSCARS)
 - Inter-domain Circuit Provisioning Example
 - Network Description Language
 - Using Static Tunnels
- References

Goals of Network Stitching

- Establishment of paths across multiple domains
- Typically provides dedicated network resources/QoS
 - service is deterministic in nature
 - » guaranteed latency, jitter, maximum transmission size
 - provides a guaranteed capacity
 - » links are not oversubscribed
- May also be used to provide below IP networking connectivity
 - transmission of non-IP traffic
 - allow for experimentation with new protocols
 - HDTV transmission directly over wavelengths,
InfiniBand/FibreChannel over LAN-PHY or WAN-PHY

Our view on Deterministic Services

- Importance in experiment repeatability
 - researchers must be able to reproduce experiments under the same conditions in order to compare results
 - QoS not only at the network level, but also at the CPU level
 - » running my experiment in different parts of the “GENI cloud” may result in dramatically different results
- Network congestion should not impact experiments
- Some experiments/applications are more sensitive to network performance & require deterministic service
 - real-time or interactive applications (radio astronomy, video conferencing, remote steering)

Challenges of Network Stitching

- Multi-domain, multi-technology environment
- Differences in underlying network technologies
 - Ethernet
 - » 802.3 (untagged Ethernet)
 - » 802.1Q (single VLAN ID)
 - » 802.1ad (provider bridges: Q-in-Q or VLAN stacking - inner/outer VLAN IDs)
 - » 802.1ah (provider backbone bridges: MAC-in-MAC, complete recursion of headers, tunneling of 802.3, 802.1Q or 802.1ad)
 - SONET/SDH
 - » virtual versus contiguous concatenation
 - all-optical (OOO) DWDM
 - » wavelength spacing/alignment (e.g., 50GHz versus 100GHz)


Challenges of Network Stitching

- Differences in underlying network technologies
 - Layer 2 MPLS Virtual Leased Line (VLL)
 - » Martini draft, VLAN circuit cross-connect
 - » MPLS LSP stitched to an Ethernet VLAN on the edges
 - IP MPLS w/ QoS
 - Premium IP
 - » Traditional IP /w QoS, mark packets with DiffServ Code Point
- Automatic versus Manual provisioning
 - Some domains support automatic establishment of network paths through a control plane
 - Manual configuration at the equipment level is still required in many cases

End-to-End Slices Across Network Aggregates


- Ontology adopted by the DCN Software Suite
 - Domains
 - Nodes
 - Ports
 - Links
 - » captures technology-specific attributes (e.g., available VLAN IDs for Ethernet, available timeslots for SONET/SDH)
- Based on OGF Working Groups:
 - Network Measurements Working Group (NM-WG)
 - » <https://forge.gridforum.org/sf/projects/nm-wg>
 - Network Mark-up Language Working Group (NML-WG)
 - » <https://forge.gridforum.org/sf/projects/nml-wg>

Domain



URN identifying a particular domain
urn:ogf:network:domain=blue.pod.lan

Nodes




3 URNs identifying switching nodes

urn:ogf:network:domain=blue.pod.lan:node=switch1

urn:ogf:network:domain=blue.pod.lan:node=switch2

urn:ogf:network:domain=blue.pod.lan:node=switch3


Ports



2 URNs identifying physical switch ports on node switch2
urn:ogf:network:domain=blue.pod.lan:node=switch2:port=4
urn:ogf:network:domain=blue.pod.lan:node=switch2:port=5

Link

Note: One physical port *may* contain multiple logical links



2 URNs identifying link endpoints for a single link

urn:ogf:network:domain=blue.pod.lan:node=switch1:port=5:link=11.2.3.1

urn:ogf:network:domain=blue.pod.lan:node=switch3:port=5:link=11.2.3.2

XML Representation

```
<domain id="urn:ogf:network:domain=blue.pod.lan">
  <node id="urn:ogf:network:domain=blue.pod.lan:node=switch3">
    <address>192.168.2.4</address>
    <port id="urn:ogf:network:domain=blue.pod.lan:node=switch3:port=3">...
    <port id="urn:ogf:network:domain=blue.pod.lan:node=switch3:port=5">...
    <port id="urn:ogf:network:domain=blue.pod.lan:node=switch3:port=7">
      <capacity>1000000000</capacity>
      [...]
      <link id="urn:ogf:network:domain=blue.pod.lan:node=switch3:port=7:link=*">
        <remoteLinkId>urn:ogf:network:domain=red.pod.lan:node=switch1:port=6:link=*</remoteLinkId>
        <capacity>1000000000</capacity>
        [...]
        <SwitchingCapabilityDescriptors>
          <switchingcapType>l2sc</switchingcapType>
          <encodingType>ethernet</encodingType>
          <switchingCapabilitySpecficInfo>
            <interfaceMTU>9000</interfaceMTU>
            <vlanRangeAvailability>0,100-200</vlanRangeAvailability>
          </switchingCapabilitySpecficInfo>
        </SwitchingCapabilityDescriptors>
      </link>
    </port>
  </node>
</domain>
```

details hidden for other ports on this node

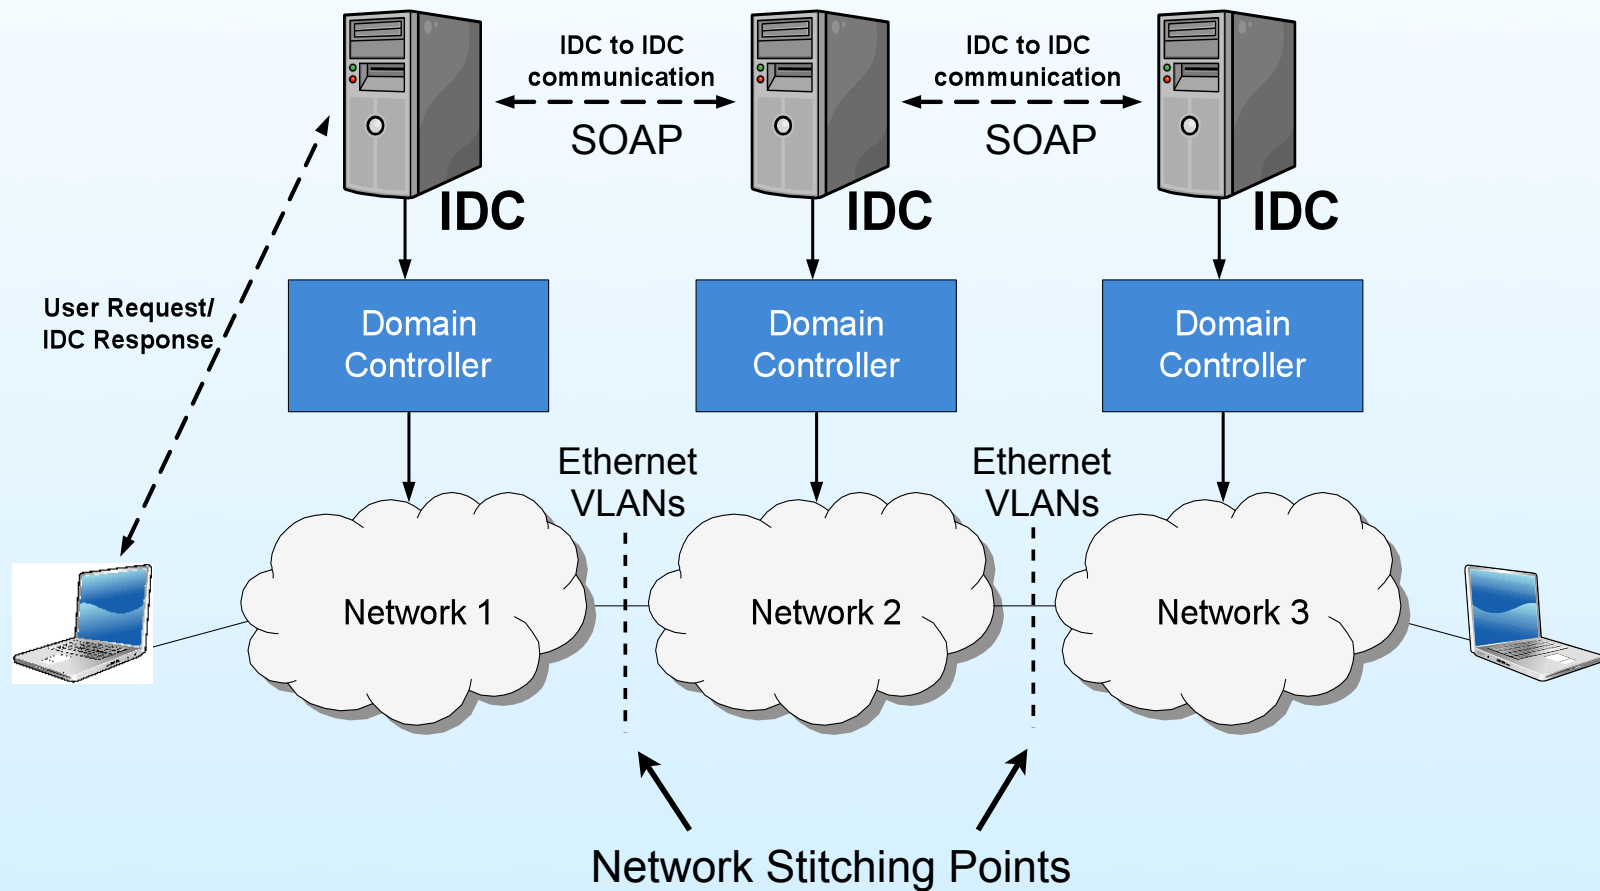
inter-domain link (blue to red)

layer 2 switching, Ethernet encoding

technology specific info for Ethernet:

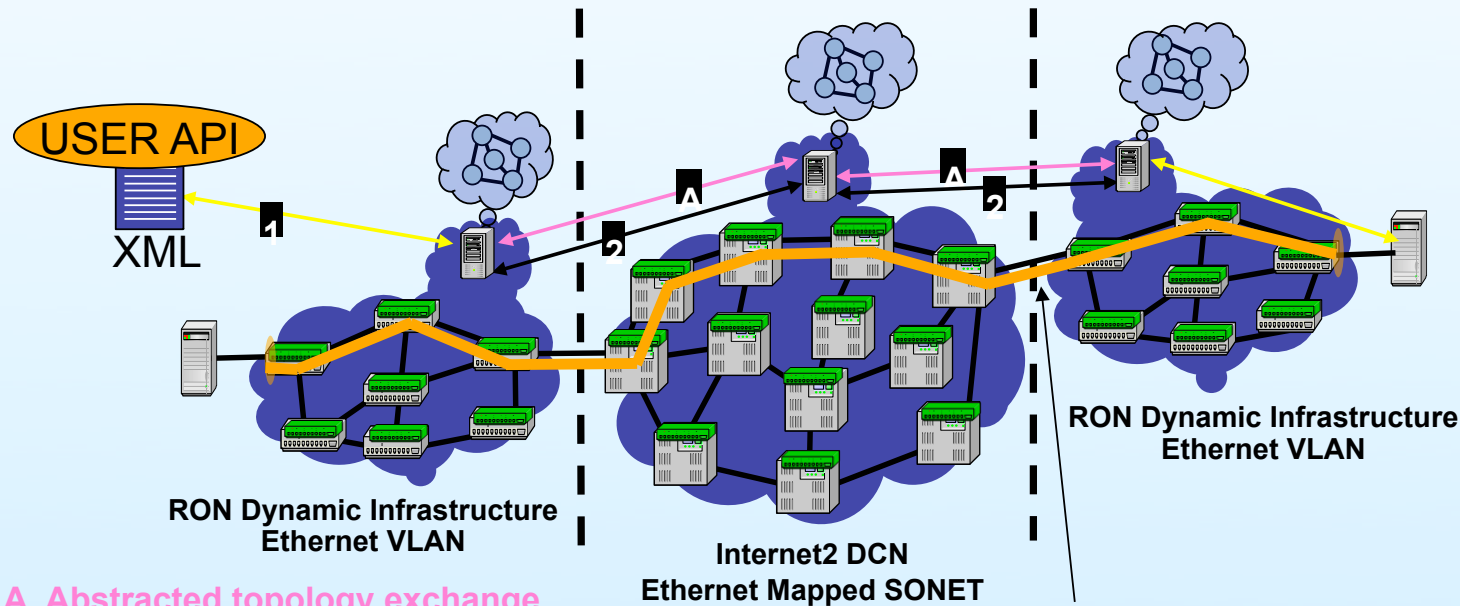
- interface MTU
- range of allowed VLANs on this link

Inter-Domain Circuit Provisioning



Dynamic Network Services Inter-Domain

- No difference from a client (user) perspective for InterDomain vs IntraDomain



A. Abstracted topology exchange

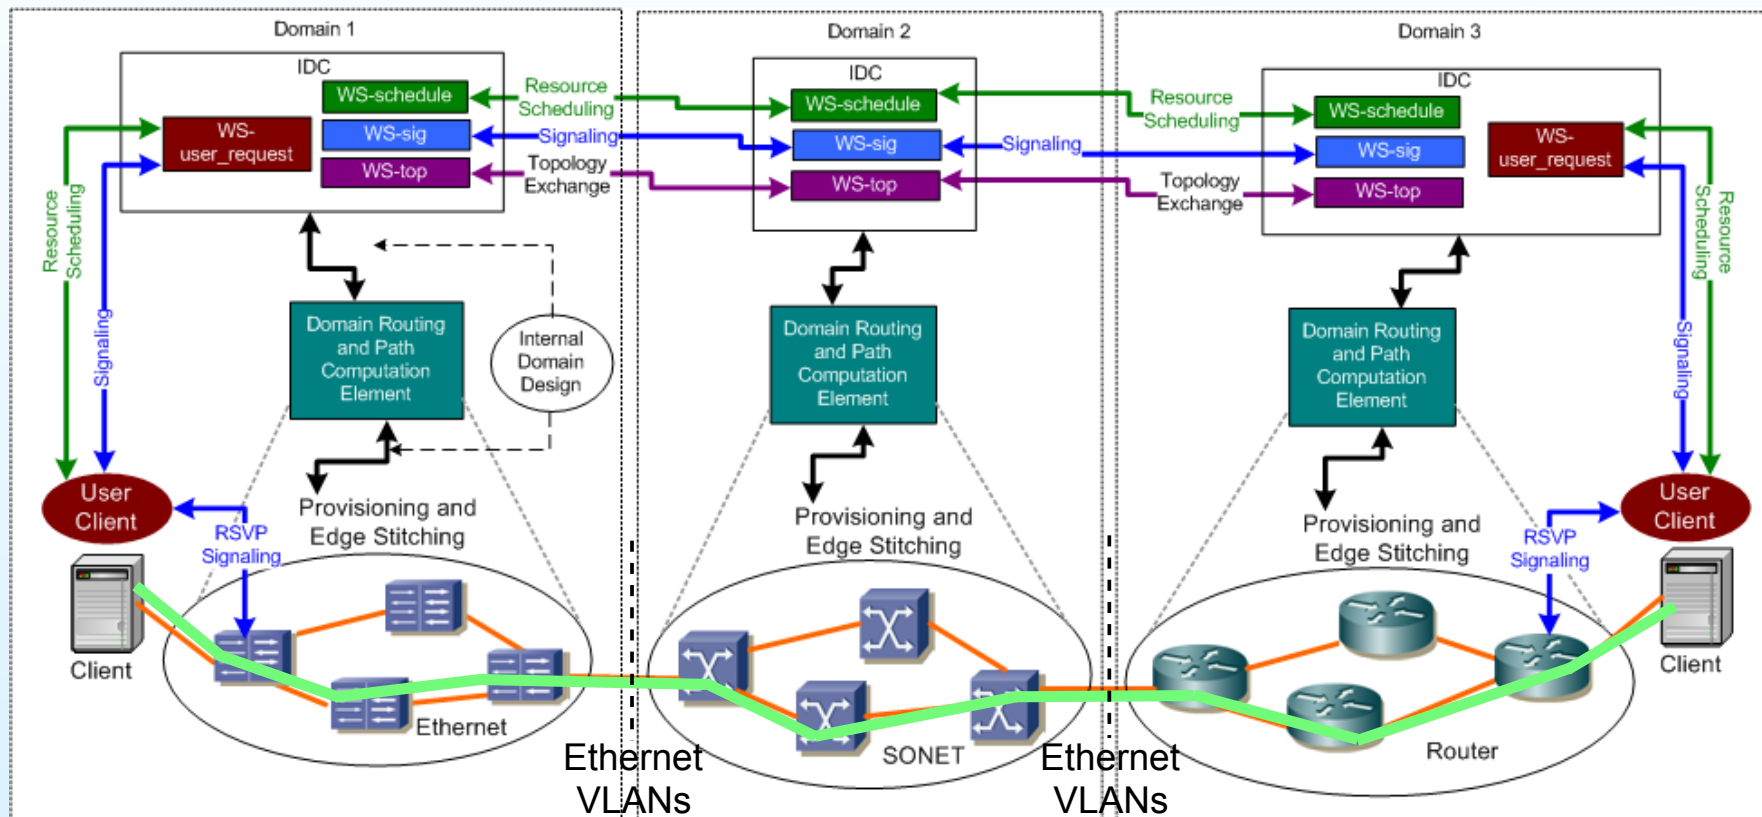
1. Client Service Request

2. Resource Scheduling

5. Service Instantiation (as a result of Signaling)

**Multi-Domain Dynamically
Provisioned Circuit**


IDC - Web Service Based Definition




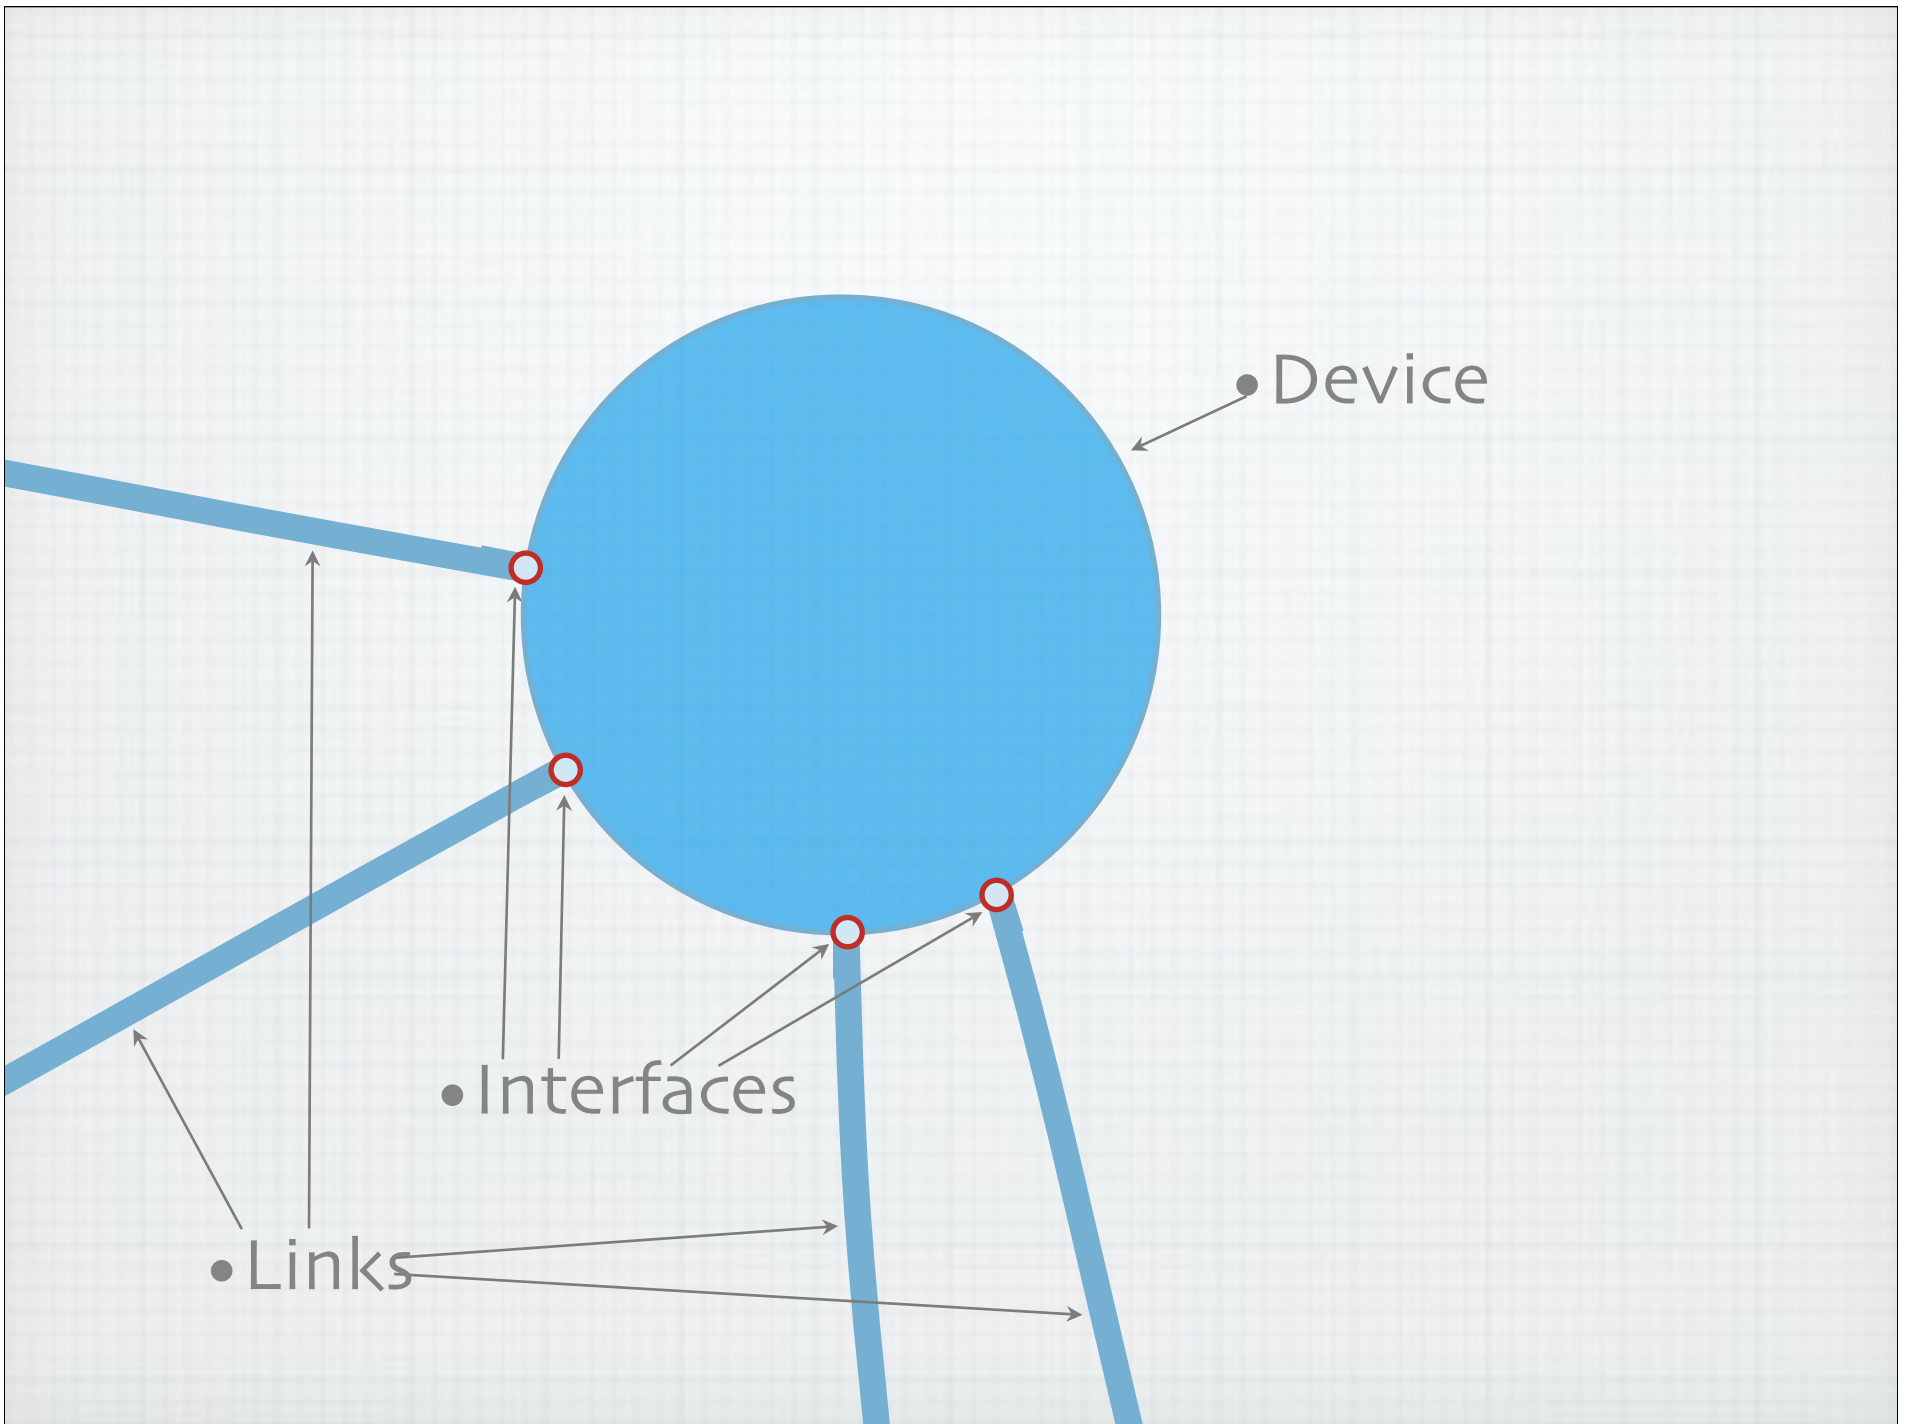
- Four Primary Web Services Areas:
 - Topology Exchange, Resource Scheduling, Signaling, User Request

Network Description Language

(slides courtesy of NML-WG at <https://forge.gridforum.org/sf/go/doc15666?nav=1>)







- Switching matrix

- A device switches data based on:


- The source interface
- One or more labels

- Example label types:

- Ethernet VLAN
- SONET STS Channel
- Wavelength (λ)

- For example, all data from channel 31 of interface 2 is forwarded to channel 28 of interface 4.

- Channels



References

- Report on Testing of Technology Stitching (DJ3.5.3)
 - <http://www.geant2.net/server/show/nav.00d00b002001003>
- Multilayer NDL
 - <https://forge.gridforum.org/sf/go/doc15666?nav=1>
- Hybrid Multi-Layer Network Control
 - Hybrid Network Dataplane Test Report
 - <http://hybrid.east.isi.edu/twiki/bin/view/HybridMLN/DataPlane>

Thanks!

- Questions or comments are welcome
- These slides will be posted at:
 - <http://geni.maxgigapop.net>
 - click on Publications
 - scroll to GEC5