

ESnet On-Demand Secure Circuits and Advance Reservation System (OSCARs)

QUILT

May 15-16, 2006

**Chin Guok (chin@es.net)
ESnet Network Engineer**

**David Robertson (dwrobertson@lbl.gov)
DSD Computer Software Engineer**

Lawrence Berkeley National Laboratory

Outline

- **Requirements for Virtual Circuit Services**
- **Virtual Circuit Characteristics**
- **Virtual Circuit Service Functional Requirements**
- **Environment of Science is Inherently Multi-Domain**
 - **Inter-Domain Virtual Circuit Requirements**
- **OSCARS: Guaranteed Bandwidth VC Service for SC Science**
- **ESnet Virtual Circuit Service Roadmap**

➤ Requirements for Virtual Circuit Services

- Identified as one of the two most important new network service by the 2002 science requirements workshop (#50-1) (the other being end-to-end performance monitoring)
- Intended to address specific science needs

➤ Virtual Circuit Characteristics

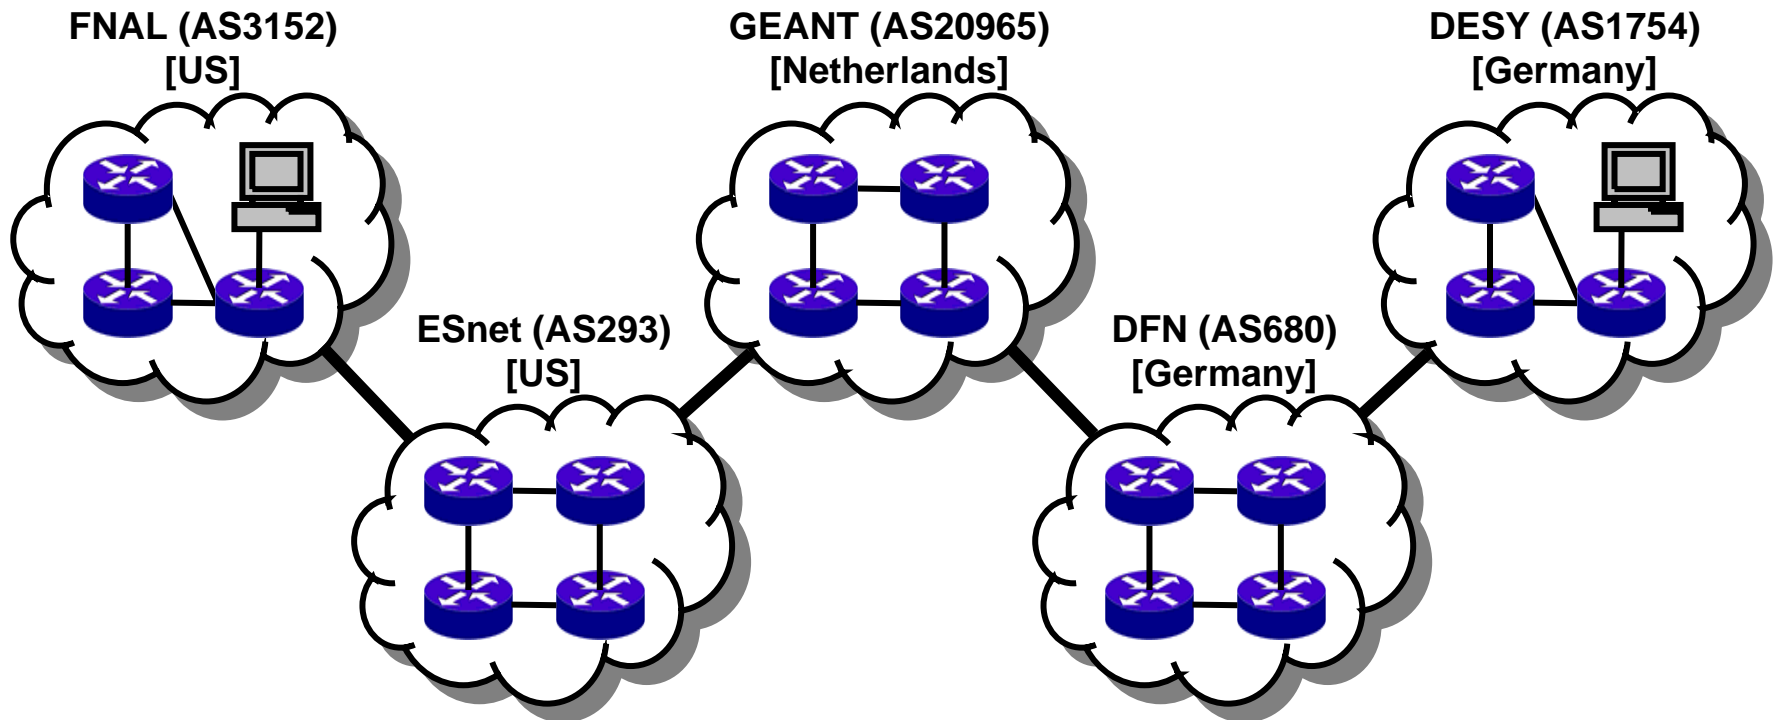
- **Traffic isolation and traffic engineering**
 - Provides for high-performance, non-standard transport mechanisms that cannot co-exist with commodity TCP-based transport
 - Enables the engineering of explicit paths to meet specific requirements
 - e.g. bypass congested links, using lower bandwidth, lower latency paths
- **Guaranteed bandwidth [Quality of Service (QoS)]**
 - Addresses deadline scheduling
 - Where fixed amounts of data have to reach sites on a fixed schedule, so that the processing does not fall far enough behind that it could never catch up – very important for experiment data analysis
- **Reduces cost of handling high bandwidth data flows**
 - Highly capable routers are not necessary when every packet goes to the same place
 - Use lower cost (factor of 5x) switches to relatively route the packets
- **End-to-end connections are required between Labs and collaborator institutions**

➤ Virtual Circuit Service Functional Requirements

- **Support user/application VC reservation requests**
 - Information required in the request include
 - Source and destination of the VC
 - Bandwidth, start time, and duration of the VC
 - Traffic characteristics (e.g. flow specs) to identify traffic designated for the VC
- **Manage allocations of scarce, shared resources**
 - Authentication to prevent unauthorized access to this service
 - Authorization to enforce policy on reservation/provisioning
 - Gathering of usage data for accounting
- **Provide circuit setup and teardown mechanisms**
 - Widely adopted and standard protocols (such as MPLS and GMPLS) are well understood within a single domain
 - Cross domain interoperability is the subject of ongoing, collaborative development
- **Enable the claiming of reservations**
 - Traffic destined for the VC must be differentiated from “regular” traffic
- **Enforce usage limits**
 - Per VC admission control is needed to police usage, which in turn facilitates guaranteed bandwidth
 - Consistent per-hop QoS throughout the network facilitates transport predictability

➤ Environment of Science is Inherently Multi-Domain

- End points will be at independent institutions – campuses or research institutes - that are served by ESnet, Abilene, GÉANT, and their regional networks
 - Complex inter-domain issues – typical circuit will involve five or more domains
 - For example, a connection between FNAL and DESY involves five domains, traverses four countries, and crosses seven time zones



➤ Inter-Domain Virtual Circuit Requirements

- **In order for end-to-end VCs to work, domains must agree and coordinate**
 - Management Plane
 - Control Plane
 - Data Plane

Inter-Domain VC Management Plane Requirements

- **Management Plane**

- Authentication and authorization between domains must be compatible
- Bi-lateral agreements between domains to enforce acceptable use policies

Inter-Domain VC Control Plane Requirements

- **Control Plane can be implemented in two ways**
 - Single VC signaled end-to-end
 - Domains must have a compatible control plane mechanism to setup and teardown end-to-end VCs
 - Authentication and authorization of control messages required to prevent unauthorized provisioning
 - This is currently lacking in vendor implementations and so this approach cannot be used, even between very closely cooperating networks
 - Multiple VCs stitched together at clearly defined places (e.g. domain borders)
 - Domains must have a mechanism for neighboring domains to request VCs
 - Domains must have pre-knowledge or the ability to locate these reservation mechanisms in neighboring domains
 - End-to-end path discovery is required to know where to stitch the VCs together

Inter-Domain VC Data Plane Requirements

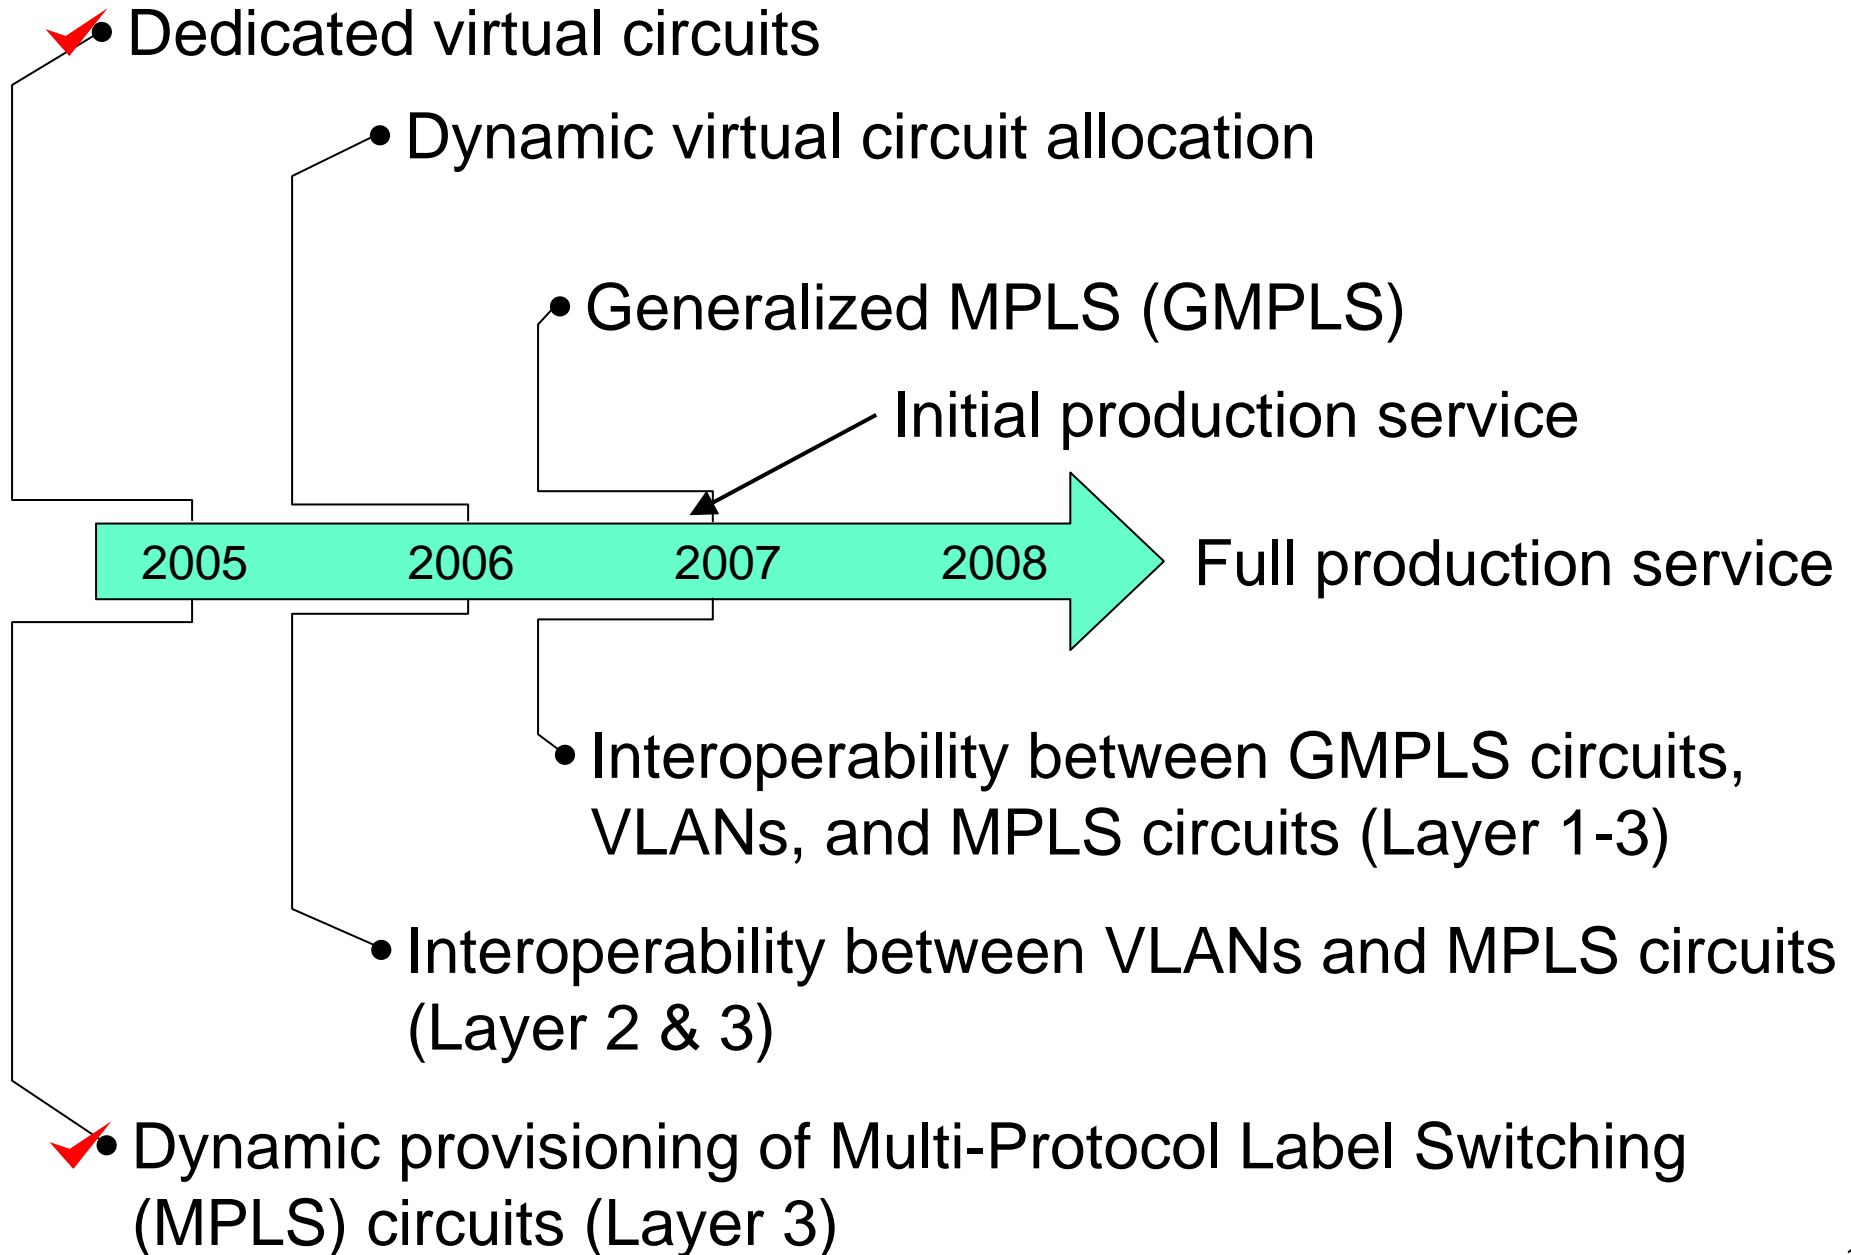
- **Data Plane**

- Transport handoff between domains must be the same (e.g. IP, VLANs, PoS)
- Bandwidth guarantee requires consistent end-to-end QoS and admission control

➤ OSCARS: Guaranteed Bandwidth VC Service For SC Science

- **ESnet On-demand Secured Circuits and Advanced Reservation System (OSCARS)**
(#50-2)
- **To ensure compatibility, the design and implementation is done in collaboration with the other major science R&E networks and end sites**
 - Internet2: Bandwidth Reservation for User Work (BRUW) (#50-3)
 - Development of common code base
 - GEANT: Bandwidth on Demand (GN2-JRA3), Performance and Allocated Capacity for End-users (SA3-PACE) and Advance Multi-domain Provisioning System (AMPS) (#50-4)
Extends to NRENs
 - BNL: TeraPaths - A QoS Enabled Collaborative Data Sharing Infrastructure for Peta-scale Computing Research (#50-5)
 - GA: Network Quality of Service for Magnetic Fusion Research (#50-6)
 - SLAC: Internet End-to-end Performance Monitoring (IEPM) (#50-7)
 - USN: Experimental Ultra-Scale Network Testbed for Large-Scale Science (#50-8)
- **In its current phase this effort is being funded as a research project by the Office of Science, Mathematical, Information, and Computational Sciences (MICS) Network R&D Program**
- **A prototype service has been deployed as a proof of concept**
 - To date more than 20 accounts have been created for beta users, collaborators, and developers
 - More than 100 reservation requests have been processed

➤ ESnet Virtual Circuit Service Roadmap



Footnotes

- 5b-1 Report of the High Performance Network Planning Workshop
http://www.es.net/pub/esnet-doc/2-3high-performance_networks.pdf
- 5b-2 ESnet OSCARS webpage: <http://www.es.net/oscars>
- 5b-3 Internet2 BRUW Project: <http://discvenue.internet2.edu/wordpress>
- 5b-4 GEANT PACE Project: <http://pace.geant2.net>
- 5b-5 BNL TeraPaths Project: <http://www.atlasgrid.bnl.gov/terapaths>
- 5b-6 General Atomics QoS Project: <http://www.fusiongrid.org/network>
- 5b-7 SLAC IEPM Project: <http://www-iepm.slac.stanford.edu>
- 5b-8 UltraScienceNet Testbed: <http://www.usn.ornl.gov>