

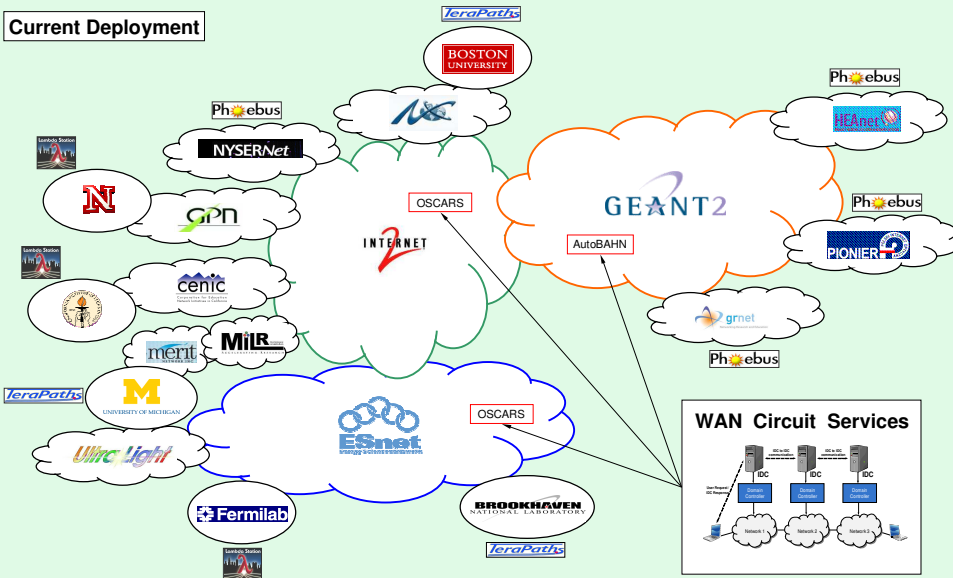
# A Collaborative Network Middleware Project by



, **TeraPaths** , and

**Phoebus**

## Current Deployment



The TeraPaths, Lambda Station, and Phoebus projects were funded by the Department Of Energy's (DOE) network research program to support efficient, predictable, prioritized petascale data replication in modern high-speed networks, directly address the "last-mile" problem between local computing resources and WAN paths, and provide interfaces to modern, high performance hybrid networks with a low entry barrier for users. Within the framework of the three projects, we successfully developed services that establish on-demand and manage true end-to-end, Quality-of-Service (QoS) aware, virtual network paths across multiple administrative network domains, select network paths and gracefully reroute traffic over these dynamic paths, and streamline traffic between packet and circuit networks using transparent gateways. These services function as "network middleware" and improve network QoS and performance for applications, playing a critical role in the effective use of emerging dynamic circuit network services. They provide interfaces to applications, such as dCache SRM, translate network service requests into network device configurations, and coordinate with each other to setup end-to-end network paths.

Building upon the success of the three projects, which target the same user community, utilize compatible technologies, and have similar goals, we work together to research and develop the next generation of network middleware. We address challenges such as cross-domain control plane signaling and interoperability, authentication and authorization, topology discovery, and dynamic status tracking. Our roadmap is to co-design network models that ensure effective inter-domain topology discovery and network utilization, utilize the perSONAR infrastructure to monitor dynamic circuit status and measure performance, enhance Grid authentication and authorization to support inter-domain trust, and integrate our joint work with the Inter-Domain Control plane efforts (IDC). The new network middleware will be deployed and fully vetted in the Large Hadron Collider data movement environment.

**www.lambdastation.org**

- LSCorrelator - synchronizes work of all services, has control functions
- LSI - unified interface for intercommunication between LS and applications, and LS to LS
- LSI Discovery service - detects new implementations, and flow-based entities at remote and local sites
- LIFESource Scheduler - activates bandwidth allocation, monitor real time usage of resources
- NETWORK CONFIGURATOR - dynamic reconfiguring of LAN and WAN

• An alternate network path selection service

• Initiates setup & tear down of dynamic circuits

• Reroutes flow-based application traffic over circuit path

• Graceful outover and fallback on path changes

• On-demand from applications (authentication & authorization)

• Deployed between Fermilab & LHC CMS Tier-2 sites at Univ. of Nebraska and Caltech

**www.terapaths.org**


- Transparently provides end-to-end (end-host-to-end-host) paths with **QoS guarantees** at the individual data flow level
- Regulates, classifies (prioritizes) traffic
- Directs traffic into dynamically or statically configured WAN paths
- Directly configures and site LAN devices: utilizes DiffServ to condition traffic and PBR to direct traffic into L2 paths
- Automatically reserves L2 paths (circuits) within ESnet SDN and I2 DCN and L3 paths (MPLS tunnels) within ESnet SDN by invoking the OSCARS API
- Bandwidth of L2/L3 paths can be shared by multiple flows with QoS guarantees for each individual data flow

**Phoebus**



<http://e2epi.internet2.edu/phoebus.html>

- Phoebus Gateways are placed at network access points and can be thought of as network "on-ramps"
- At a high level, the goal is to get network data to these gateways which can then intelligently move the data as efficiently as possible to the other edge of the network
- Phoebus will automatically utilize DCN resources as needed and as available
- At the other edge, it again appears as legacy traffic does today, thus neither client or server need to be modified

## 1st goal: interoperability



## 2nd goal: create the next generation: ESCPS (End Site Control Plane System)



CHEP 2009



**BROOKHAVEN**  
NATIONAL LABORATORY

**Fermilab**

