FERMILAB-CONF-09-727-CD

A Collaborative Network Middleware Project by Lambda Station, TeraPaths, and Phoebus

A Bobyshev [1], S Bradley[2], M Crawford[1], P DeMar[1], D Katramatos[2,4], K Shroff[2], M Swany[3] and D Yu[2]

Fermi National Accelerator Laboratory, Batavia, IL 60510, USA
{bobyshev, crawdad, demar}@fnal.gov
[2] Brookhaven National Laboratory, Upton, NY 11973, USA
{bradley, dkat, shroffk, dtyu}@bnl.gov
[3] University of Delaware, DE 19716, USA
swany@cis.udel.edu

Abstract. The TeraPaths, Lambda Station, and Phoebus projects, funded by the US Department of Energy (DOE), have successfully developed network middleware 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 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 up end-to-end network paths. The End Site Control Plane Subsystem (ESCPS) builds upon the success of the three projects by combining their individual capabilities into the next generation of network middleware. ESCPS addresses challenges such as cross-domain control plane signalling and interoperability, authentication and authorization in a Grid environment, topology discovery, and dynamic status tracking. The new network middleware will take full advantage of the perfSONAR monitoring infrastructure and the Inter-Domain Control plane efforts (IDC) and will be deployed and fully vetted in the Large Hadron Collider data movement environment.

---

[14] Author to whom any correspondence should be addressed.
2

3

1. Introduction

The End Site Control Plane Subsystem (ESCPS) is a project to research, develop, and deploy an end site component of the federated control plane capability that provisions dynamic circuits across multiple administrative network domains with advanced tools and services for diagnosing faults and monitoring the performance of dynamic circuits. The objective is to make on-demand, dynamic circuit services a truly end-to-end construct, rather than merely from one site border to another. ESPCS will incorporate or adapt existing and emerging control plane technologies, such as DiffServ, MPLS, and GMPLS, for control and configuration of local site network infrastructures and combine these with wide area dynamic circuits. Authentication and authorization (AA) capabilities will be integrated into the evolving federated cross-domain control plane AA infrastructure emerging within the research & education wide-area network providers (ESnet, Internet2, GEANT). Reservation specifications and parameters, as well as control plane signalling protocols, will be supported on an Inter-Domain Controller (IDC) component of ESCPS that coordinates with wide-area IDC elements to establish, manage, and tear down end-to-end circuits. The ESCPS will provide critical end-to-end services the wide-area backbone control plane infrastructure cannot provide; 1) continuous operational status monitoring; 2) performance analysis of data traffic traversing the circuit which checks for expected performance characteristics, and potentially initiates reconfiguration actions if necessary; and finally 3) troubleshooting capabilities that can localize, identify, and provide domain-specific notification on problems affecting the connection status or performance.

2. Background

Emerging dynamic circuit services within the research and education (R&E) backbone network providers, such as ESnet and Internet2, present new opportunities for the advancement of wide-area distributed computing. A primary motivation for the development of these services is high impact data movement by large-scale science collaborations. The Large Hadron Collider (LHC) experiments, coming online in 2009, represent the vanguard of next-generation science collaborations that are expected to be based on truly distributed computing models. LHC-related traffic already represents a sizeable majority of the traffic carried across the ESnet and Internet2 backbones. Other large-scale science collaborations, including climate modelling and fusion research programs are projected to generate similar levels of data traffic on global scales. The capability to dynamically configure data paths with bandwidth guarantees and traffic isolation offer these international collaborations a higher level of assurance that network performance will adequately meet the needs of their globally distributed computing systems. A corollary impact of circuit-based network services is the capability to separate large-scale science data traffic from general internet traffic, preventing the latter's network-sensitive applications from incurring congestion-induced performance degradation, due to high impact data movement of the former. One should even extrapolate that those network-sensitive applications of a general nature, such as high-definition videoconferences, may eventually become users of circuit-based network services themselves. In short, the development and advancement of circuit-based services has become a key strategic direction in the evolution of network technology within the R&E community.

2.1. Problem Description

Dynamic circuits represent the only scalable solution to large-scale deployment of circuit-based network services on a global level. The ability to reserve, setup and tear down a data circuit quickly and in an automated manner across multiple, administratively distinct network domains is a basic requirement for a widely deployable circuit-based network service. Research efforts in developing a cross-domain control plane infrastructure for dynamic circuit services have made considerable progress over the past several years. The integration of ESnet's On-Demand Secure Circuits and Advanced Reservation Service (OSCARS) circuit reservation project with the Internet2's Dynamic Resource Allocation via GMPLS Optical Networks (DRAGON) control plane development project has led to the deployment of a dynamic circuit service capability, called DCN/OSCARS, across the Internet2 and ESnet backbones. GEANT, the trans-European R&E backbone network is developing a comparable dynamic circuit capability, named AutoBAHN. Close coordination of the DCN/OSCARS and AutoBAHN development efforts are expected to provide a common dynamic circuit service with integrated control plane capability between the US and European R&E backbone networks. However, these efforts have been limited in scope to the domains of the wide-area network providers. The control plane solution essentially stops at the perimeter of the end sites involved. Put another way, the data circuit is limited to being a site-to-site construct, not a true end-to-end construct between data source and data sink. As large-scale collaborative science evolves toward petascale computing and terabit networks, limiting the scope of the control plane to the wide-area service provider infrastructure will likely be an inadequate solution.

There are several other limitations with the existing control plane model. First, there is no end-to-end circuit monitoring within the circuit service. While each wide-area domain may be capable of monitoring its segment of the path, there is no end-to-end monitoring within the control plane to insure the circuit is functional, or performing within expectations. The end sites may well be providing their own end-to-end circuit monitoring capability, but it is disjoint from the control plane. This severely limits the responsiveness of the control plane to circuit malfunctions or problems. Second, the end sites still need to establish their own "virtual" control plane capability for setup and teardown of the circuit, in order to synchronize modification of their respective local routing configurations to utilize the circuit. Optimally, this would be part of the control plane communications, not disjoint from it.

There are several ongoing projects, such as Lambda Station (Fermilab) and TeraPaths (BNL), which facilitate use of DCN/OSCARS dynamic circuits by end sites, adapted for their specific environments. However, each product was developed to interface its end site to DCN/OSCARS services, not to extend the circuit and its control plane into the site. Currently, there is no tool a site might use to create a true end-to-end circuit that facilitates unified signalling across site, wide area, and DCN networks

2.2. Proposed Solution

The End Site Control Plane Subsystem (ESCPS), will be the end-site component of the federated control plane capability, that provisions dynamic circuits across multiple administrative network domains. The overall objective of the project is to make on-demand, dynamic circuit services a truly end-to-end construct, rather than limited to the boundaries of wide-area service providers. Figure 1 illustrates the concept. An ESCPS server deployed at a site is an integrated part of the federated control plane to establish and manage dynamic circuits. The ESCPS server performs the local network control functions required to establish the circuit. Finally, the ESCPS server instigates creation of the circuit upon receipt of a request from a user, application, or other network trigger, assuming the request meets authentication and authorization constraints.

ESCPS will incorporate or adapt, as appropriate, existing and emerging control plane technologies, such as MPLS, and GMPLS, for control and configuration of local site network infrastructures, integrating these with wide area dynamic circuit control and configurations. Authentication and authorization (AA) capabilities would conform to the evolving federated cross-domain control plane AA infrastructure emerging within the research & education wide-area network providers (ESnet, Internet2, GEANT). Reservation specifications and parameters, as well as control plane signalling protocols, would be supported on an Inter-Domain Controller (IDC) component of ESCPS that coordinates with wide-area IDC elements to establish, manage, and tear down end-to-end circuits. ESCPS will provide critical end-to-end services the wide-area backbone control plane infrastructure does not provide:

- Continuous operational status monitoring to verify a circuit's end-to-end connectivity, a function not currently in existence within the emerging research & education dynamic circuit infrastructure
- Performance analysis of data traffic traversing the circuit for expected performance, potentially initiating reconfiguration actions if it is not
- Troubleshooting capabilities that could localize, identify, and provide domain-specific notification on problems affecting connection status or performance

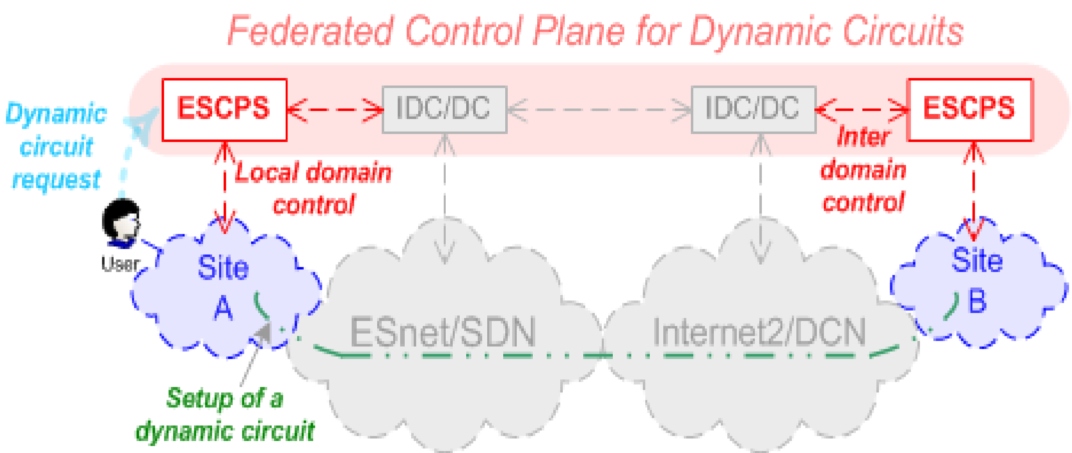

Figure 1. ESCPS concept

2.3. Significance

ESCPS will directly supports ongoing development efforts in four areas expected to become cornerstones of R&E network technology, as it evolves toward support of petascale science, based on terabit networks:

- Advancing the deployment of dynamic circuit services: ESCPS will provide the necessary dynamic circuit monitoring capabilities needed for a robust, production-quality service. It will represent a single, unifying package for local site support of dynamic circuit services.
- Enhancing evolution of dynamic circuit technology: Convergence on a common, truly end-to-end control plane would provide the potential to develop advanced control plane features that increase reliability and reconfiguration agility.
- Encouraging deployment of monitoring infrastructure: perfSONAR monitoring infrastructure deployment has been slow, due to lack of applications that use it. ESCPS will create strong incentive for sites to deploy perfSONAR infrastructure.
- Encourage standardization on emerging Grid standards. Our product will adhere to Open Grid Forum (OFG) interface and markup language standards, effectively pushing those standards into applications that make use of our services.

3. ESCPS Design & Architecture

End Site Control Plane Subsystem is based on a three-layer reference model and consists of control plane, service management and interface layers. Each layer has a well-defined set of the functionalities and interfaces to other layers and external entities. All components will be structured as modules that can conceivably run as independent Web Services. The Client API exposes the services of the system to applications, middleware, users and network traffic engineering agents. The Request Manager accepts requests and verifies that incoming requests are allowed based on interactions with the Authentication, Authorization and Accounting (AAA) service or services. ESCPS will rely on external Authentication (AuthN) and will interface with external Authorization (AuthZ) services to provide security attributions and policy. In addition, ESCPS will add an additional Policy Control service to express local policies for final Authorization. The Resource Manager (RESMGR) brokers and maintains state about ongoing resource utilization. The RESMGR will negotiate with the End-to-End Path Services module that will arrange the end-to-end path using back-ends to provision the various components from the host, the local domain and the wide-area, via the Inter-Domain Controller (IDC). The initiating ESCPS instance may also communicate with a remote ESCPS instance at the destination site.

3.1. Architecture

The ESCPS architecture is depicted in Figure 2. It is based on three layers reference model. The layers, control plane, service management and interface have the well-defined sets of functionalities and interfacing each other and with external entities via the Web-based Services.

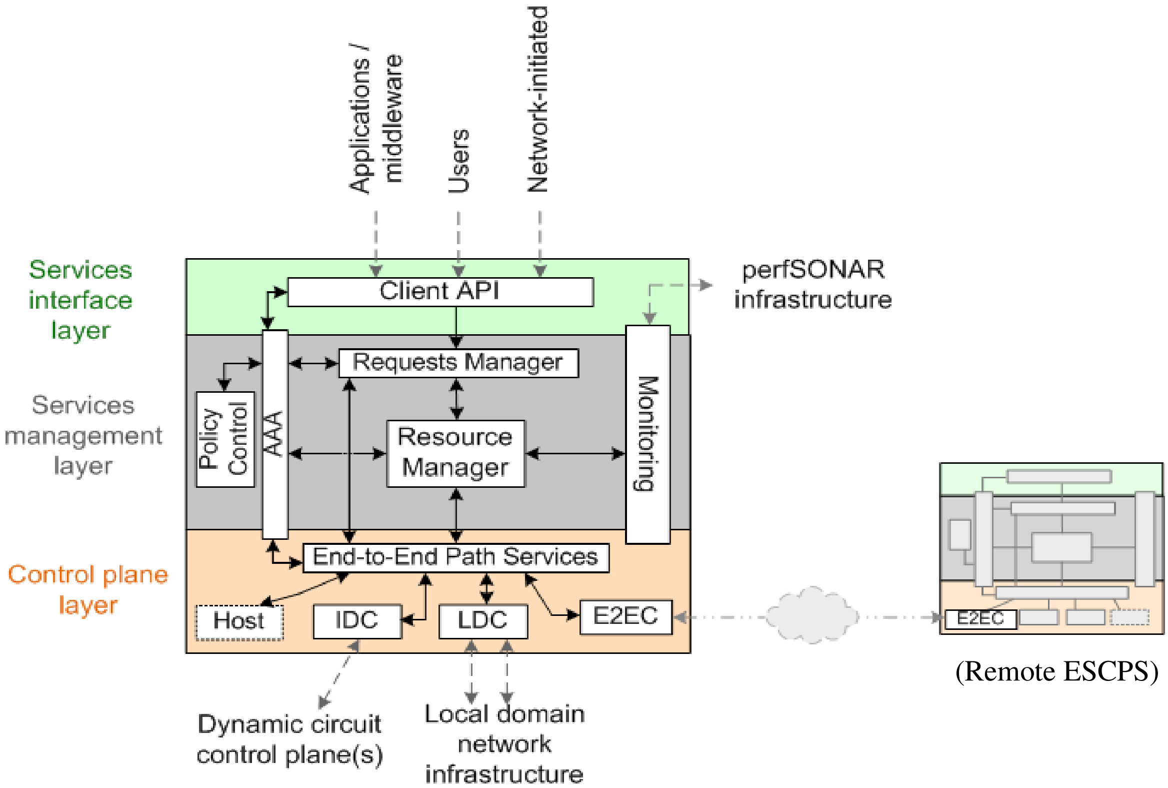

Figure 2. ESCPS Architecture

3.1.1. Control plane layer. E2EPS as a set of "proxy" module with a uniform interface for both WAN or end-site services. Each "proxy" module uses different communication protocols to interacts with the targeted host, WAN service, such as (DOE ESnet OSCARS, and Internet 2 DCN services), and the flexibility of the design allows an E2EPS to interface with possibly any kind of domain controller that offers at least basic path setup functionality; it is similar to the way operating systems control hardware through device drivers. Furthermore, it is possible to follow any end-to-end setup coordination model. ESCPS follows a hybrid star/daisy chain coordination model where the initiating end-site first coordinates with the target site and then indirectly sets up a WAN path by contacting its primary WAN provider and relying on that provider's domain to coordinate, if necessary, with other WAN domains along the desired route. This approach does not require an end site to set up a route by individually contacting WAN domain controllers along a WAN chain (star model), which would, in turn, require end-sites to have detailed knowledge of the network so as to know which controllers need to be contacted, what capabilities each one has, etc. The star model is centralized and significantly increases the complexity of setting up a path. On the other hand, the daisy chain model requires all participants (end-sites and WAN domains) to use a common communication protocol that allows full functionality of all basic operations of every participant. Such a protocol does not yet exist. The hybrid coordination model was thus adopted as the most feasible.

3.1.2. Service Management Layer. The Service Management layer has two main components, Requests Manager and Resource Managers. Applications middleware, users (via a web-based *portal), other control services submit requests to ESCPS via Request Manager (REQMGR). In response, REQMGR returns an identifier assigned to that request. When time specified in requests arrived REQMGR starts processing all current local requests simultaneously as well as initiates the corresponding procedures at remote end by using services of the lower layer.*

The Resource Manager (RESMGR) performs functions of a local site's bandwidth broker and initiates communication protocol to RESMGR at a remote end to perform corresponding steps. The limitation in the available bandwidth and additional overhead of circuits to reserve bandwidth requires sophisticated techniques to maximize utilization efficacy and scalability in terabit networks. We experienced significant overhead of circuit creation and tear-down in federated network environment to selected prioritized application flows. This Layer will also address the issues of scalability and bandwidth utilization in the LAN advanced reservation, allocation, and scheduling with an innovative strategy of flow grouping and aggregation, i.e. flows with matching source and destination need to be grouped together and forwarded through common circuits, configured so that they accommodate the aggregate bandwidth of the grouped flows. Grouping together individual data flows or flow groups with common source and destination and forwarding them to a common circuit with enough total bandwidth and duration to accommodate all flows can drastically reduce the number of circuits that are needed between a pair of end sites simultaneously and increase the availability of the dedicated paths. End sites still reserve resources for individual flows, however multiple end site reservations can be accommodated by a single circuit reservation as long as the aggregate duration and bandwidth can be determined. The level of reservation consolidation (or unification) needs to be controlled by suitable criteria to minimize waste of resources. We will consider several optimization and consolidation techniques for circuit reservations.

The standard approach used by most Advanced Resources Reservation Systems deployed in DC networks today is to use information provided in submitted requests. This information, however, might be not very accurate due to several reasons such as difficulty to estimate actual needs for the bandwidth, a difference between needed and achievable bandwidth. Due to these and other reasons a user may always wish to allocate a maximum allowed bandwidth but, in fact, not be able to use it. As a result expensive resources will not be used and competing requests for same time slot will be rejected. To avoid inefficient resource utilization RESMGR can use information from the Monitoring Module to adjust allocated and actual utilization of resources. The RESMGR will use adaptive algorithms to achieve efficiency of resource utilization based on analysis of the past utilization, the results of active monitoring and forecast models.

3.1.3. Interface Layer, Policy control and AAA. ESCPS provides two basic functions to the high level applications and middleware: end sites status including logical topological information of LAN, existing end to end paths reservations, and circuit operations, such as reservation, modification and cancellation. The services provided by ESCPS are exported via the web services containers. We plan to provide three mechanisms for client applications to use the end-to-end paths: web and command line interfaces for manual path reservation, and programming API and underlying libraries for application and middleware to dynamically reservation network paths. Several primitives used by clients will be implemented: reserve, commit, and cancel a path, get bandwidth information and status of a path.

We will coordinate our work on data types and Network Service API with the Network Services Interface group [10]. We will implement the application plug-ins, such as dCache SRM [12] plugin and GridFtp plugin [11], that are ESCPS clients, calls those ESCPS API to perform circuit operations. These plugins will ensure these upper layer applications and middlwares to be network-aware.

The Policy Control, Authentication, Authorization and Accounting (AAA) are spread across all layers. Each layer also provides monitoring information. In addition, Monitoring Module will interact with external monitoring entities such as perfSONAR to watch for the status of dynamic circuits, if available.

3.1.4. Circuit monitoring & performance analysis capabilities. We seek to provide reliable indicators that an ESCPS-created circuit is both operational and providing expected levels of performance. These are significant technical challenges. The major difficulty is in determining the status of a circuit is that it needs to be monitored from the circuit end points to be accurate. The end points are normally routers, with closed operating systems and proprietary solutions for service monitoring. Since we require an open solution, we will base our circuit monitoring development efforts on perfSONAR infrastructure and SNMP MIB objects that are commonly available. It is likely that our circuit status monitoring will involve collection and tracking of multiple MIB objects and values, which in combination or alone may indicate a circuit fault condition.

Our objective in terms of circuit performance analysis support is to be able to determine that the circuit is capable of providing the configured level of service. Our technical challenge will be to distinguish between host-based or application-based issues limiting performance versus circuit-based issues. We will investigate the feasibility of utilizing perfSONAR's emerging active performance monitoring services to help determine the circuit's performance capabilities. We will also capitalize on the Structured Performance Analysis Methodology (spelling) [ref] work done at Fermilab in approaching analysis of circuits for expected performance.

Finally, we intend to explore vendor-proprietary router solutions for service level measurement, with two objectives in mind. First, to evaluate how effective these solutions are, and whether their approach has merit for adaptation to our needs. Second, to track the course of these technologies in standards bodies, and advocate for adoption of promising ones. Initially, we will explore use of Cisco Object Tracking as service level measurement tool.

3.1.5. Fault tolerance. ESCPS through the use of existing network monitoring tools will periodically compose a current system view and compare it against earlier system views, to detect problems and initiate recovery procedures. Because ESCPS operates on behalf of the end user and has direct control only over participating end sites, it can only indirectly initiate recovery in other domains, if these domains expose such a recovery service through their domain controllers. We aim to develop and evaluate fault recovery models similar to the ones used in MPLS traffic engineering to reroute traffic in the presence of faults. In a local repair model, individual (physical) network domains should be at least able to detect faults within their network segment. Upon detection of a problem, further action depends on the capability and/or policy of the individual domain to address that problem in a way that won't affect other domains. If the domain is not able to deal with the problem in the short timescale required to avoid disruption of connections, it should at least notify other affected domains so that other courses of action can be followed, i.e., adjacent domains may attempt to bypass the problematic segment, or the *end sites may initiate an overall re-routing attempt. In a global repair model, ESCPS end site instances responsible for maintaining a virtual domain periodically check the status of the network paths composing the virtual domain using available monitoring data. Failure recovery will be triggered if the periodic check detects a problem or if a notification arrives from any other domain involved. In the first case, the ESCPS instance that detected the problem will notify all domains, including the one that has the problem and subsequently decide, also depending on other domain responses, whether to temporarily re-route traffic over a best-effort path and wait for the problem to be fixed or attempt to obtain a new circuit altogether.*

3.1.6. Troubleshooting capability. One of the enhanced services that ESCPS will provide is a set of basic troubleshooting capabilities for circuits that have been identified as malfunctioning. The current state of troubleshooting end-to-end circuits, particularly across multiple network domains, has to be classified as primitive. The rich set of IP-based troubleshooting tools that work quite across routed network paths provide little useful information across VLAN-based network paths. The problem is further complicated by the potential complexity and heterogeneity of the underlying network technologies supporting the circuit. MPLS and SONET, as well as native Ethernet, are commonly deployed to support data circuits, frequently in combination. Our approach to this problem will be twofold. First, we will investigate use of basic SNMP interface counters, such as link status, byte and discard counters, along the circuit traversal path to detect potential problem locations. We will make use of the emerging perfSONAR [5] monitoring infrastructure for the collection of that data. We will attempt to capitalize on the work of DANTE in setting up the E2ECU [13] to monitor the circuit-based LHC Optical Private Network, although that effort involved static, not dynamic circuits. Our second approach will be to investigate Ethernet connectivity fault management developments under the IEEE802.1ag [15] umbrella, for applicability to hybrid environment of mixed network technologies.

4. Relevant Work & Research:

ESCPS project uses experience, expertise and technical solutions obtained by the following projects.

4.1. TeraPaths.

The TeraPaths project at BNL has developed software with the capability of dynamically establishing virtual end-to-end network paths with guaranteed QoS and dedicating them to specific data flows. TeraPaths achieves this by directly configuring the networking hardware of a site's LAN and by interfacing with other domain control software systems, notably OSCARS/IDC, to indirectly configure the hardware of network domains along a requested end-to-end path. Within LANs, TeraPaths utilizes DiffServ-based QoS to protect and regulate individual flows. To carry these already conditioned flows through WANs, TeraPaths reserves paths in the form of dynamic circuits (Layer 2) and/or MPLS tunnels (Layer 3).

4.2. Lambda Station.

The Lambda Station project has been aimed to enable dynamic allocation of alternate network paths for traffic of production SciDAC applications and to forward designated flow across LAN, it negotiates with reservation and provisioning systems of WAN control planes such as DCN/OSCARS, demand tunnels, or dynamic circuit networks. It creates End-To-End circuit between single hosts, computer farms or networks with predictable performance characteristics, preserving QoS if supported in LAN and WAN and tied security policy allowing only specific traffic to be forwarded or received through created path. Lambda Station project also explored Network Awareness capabilities.

4.3. PerfSONAR, Phoebus

The perfSONAR-PS is a set of independent software services that implement the perfSONAR protocols for network performance monitoring. One of several objectives of the proposed ESCPS project is to incorporate perfSonar monitoring capabilities into control plane fault tolerance and recovery protocols of ESCPS.

4.4. DRAGON, OSCARs, and Internet2 DCN

We are aware of other relevant research and development projects such as Internet2's DRAGON, ESNet's OSCARs, GEANT2 AutoBAHN. ESnet and Internet2 have deployed Dynamic Circuit Networks (DCN) which are interconnected at multiple POPs at locations through the United States. In addition to ESnet and Internet2, GEANT's AutoBAHN test network is participating in developing and deploying DCN, with connections to United States and to NRENs in Europe. The Interdomain Control (IDC) Protocol that is used to support interoperation between networks has been developed by these three entities. The IDC protocol is based largely on the On-demand Secure Circuit Advance Reservation System (OSCARS) service [1], which is a project, initiated by ESnet and with collaboration with Internet2 for the past few years. OSCARS dynamically provisions secure guaranteed bandwidth circuits within its network and with other networks that use the IDC protocol. OSCARS initially provided guaranteed bandwidth circuits within ESnet in the form of MPLS tunnels (layer 3). Through the collaboration between ESnet and Internet2, the system evolved into a more general Inter-Domain Controller (IDC) still provides MPLS tunnels within ESnet, but that can now also provide guaranteed bandwidth layer 2 circuits within and between ESnet's Science Data Network (SDN) and Internet2's Dynamic Circuit Network (DCN). In addition, DRAGON [17] is conducting research on dynamically provisioned QoS-guaranteed connections via GMPLS.

5. Conclusions

The ESCPS aims to provide a single product with seamless, fully integrated, end-to-end control plane capability that provisions dynamic circuits across multiple administrative network domains. It builds upon the previous work done in the area of dynamic local path reconfiguration by the Lambda station and TeraPaths projects. The ESCPS will also provide end-to-end circuit monitoring functionality designed around analysis of monitoring data collected from perfSONAR monitoring infrastructure.

The ESCPS package will provide facilities a simple and manageable way to incorporate their campus networks into the DCN/OSCARS federated control plane and the monitoring and analysis functionality of ESCPS will provide means of evaluating performance of the dynamic circuits.

The ESCPS's control plane capacity to extend a circuit end-point into the campus network infrastructure will enable facilities to extend terabit scale WAN network infrastructure into their campus network at layer2 or even layer1 making it a significantly less costly option to that of deploying layer3 terabit infrastructure.

References

[1]Guok C, Robertson D, Thompson M, Lee J, Tierney B and Johnston W October 2006 "Intra and Interdomain Circuit Provisioning Using the OSCARS Reservation System" Third International Conference on Broadband Communications, Networks, and Systems, IEEE/ICST .

[2]Internet2 Dynamic Circuits Network, http://www.internet2.edu/network/dc/

[3]Bobyshev A, Crawford M, DeMar P, Grigaliunas V, Grigoriev M, Moibenko A, Petravick D, Rechenmacher R, Newman H, Bunn J, Van Lingen F, Nae D, Ravot S, Steenberg C, Su X, Thomas M and Xia Y 2006. "Lambda Station: On-demand flow based routing for data intensive grid applications over multitopology networks", *IEEE proceedings of the Third International Conference on Broadband Communications, Networks and Systems*.

[4]Katramatos D, Gibbard B, Yu D and McKee S 2006. "TeraPaths: End-to-End Network Path QoS Configuration Using Cross-Domain Reservation Negotiation". In *Proceedings of the 3rd International Conference on Broadband Communications, Networks, and Systems* (BROADNETS 2006), San Jose, California, USA.

[5]Binczewski A, Lawenda M, Lapacz R and Trocha Sz 2007. "Application of perfSONAR architecture in support of GRID monitoring", In *Proceedings of INGRID 2007 - Instrumenting the Grid, 2nd International Workshop on Distributed Cooperative Laboratories*, S.Margherita Ligure Portofino, Italy.

[6]Phoebus, High-Performance Data Transfer for Dynamic Circuit Networks, http://e2epi.internet2.edu/phoebus.html

[7]WS-I Basic Profile Version 1.1, Final Material, 2004-08-24, Editors: K. Ballinger (Microsoft), D. Ehnebuske(IBM), et al., the Web Services-Interoperability Organization.

[8]Gommans L, Laat C D, Oudenaarde B V and Taal A 2003, "Authorization of a QoS Path Based on Generic AAA". In *Future Generation Computer Systems* Vol 19, Issue 6.

[9]Mambretti et al 2003, "The Photonic TeraStream: enabling next generation applications through intelligent optial networking at iGRID2002". In *Future Generation Computer Systems* Vol 19,Issue 6 (2003) 897-908, Elsivier.

[10]OGF Network Service Interfaces NSI-WG, http://www.ogf.org/OGF24/materials/1390/NSI_charter_01.doc

[11]Allcock B, Bester J, Bresnahan J, Chervenak A, Foster I, Kesselman C, Meder S, Nefedova V, Quesnel D and Tuecke S 2001. "Data Management and Transfer in High-Performance Computational Grid Environments". In *Parallel Computing Journal*.

[12]The distributed data storage caching system (dCache) http://www.dcache.org/

[13]E2ECU Support for LHC OPN http://www.geant2.net/alice/upload/pdf/E2ECU_LHC_OPN_Presentation.pdf

[14]Methodology: Improving Wide-Area Network Transfer Performance, https://plone3.fnal.gov/P0/WAN/netperf/methodology/

[15]IEEE, 802.1ag - Connectivity Fault Management, http://www.ieee802.org/1/pages/802.1ag.html

[16]Bandwidth on-demand (AutoBAHN), http://www.geant2.net/server/show/nav.756
[17]Lehman T, Sobieski J and Jabbari B 2006. "DRAGON: A Framework for Service Provisioning in Heterogenous Grid Networks", *Communications Magazine*, IEEE, Vol:44 Issue:3.