Software-Defined Network Exchanges (SDXs): Architecture, Services, Capabilities, and Foundation Technologies

Joe Mambretti
International Center for Advanced Internet Research (iCAIR)
Northwestern University
750 North Lake Shore Drive, Suite 600
1-312-503-0735
j-mambretti@northwestern.edu

Jim Chen
International Center for Advanced Internet Research (iCAIR)
Northwestern University
750 North Lake Shore Drive, Suite 600
1-312-503-0737
jim-chen@northwestern.edu

Fei Yeh
International Center for Advanced Internet Research (iCAIR)
Northwestern University
750 North Lake Shore Drive, Suite 600
1-312-503-0738
fyeh@northwestern.edu

ABSTRACT
Software Defined Networks (SDNs), primarily based on OpenFlow, are being deployed in single domain networks around the world. The popularity of SDNs has given rise to multiple considerations about designing, implementing, and operating Software-Defined Network Exchanges (SDXs), to enable SDNs to interconnect SDN islands and to extend SDNs across multiple domains. These goals can be accomplished only by developing new techniques that extend the single domain orientation of current SDN/OpenFlow approaches to include capabilities for multidomain control, including those for resource discovery, signaling, and dynamic provisioning. Several networking research communities have begun to investigate these concepts. Early architectural models of SDXs have been designed and implemented as prototypes. These SDXs are being used to conduct experiments and to demonstrate the potentials of SDXs.

Categories and Subject Descriptors
[**Sofware Defined Networking Exchanges (SDX)**]: Software Defined Networking (SDN), Multi-Domain SDXs, and Distributed Clouds, Programmable Networks, Experimental Network Testbeds

General Terms
Design, Experimentation, Network Prototypes, Network Testbeds, Verification.

Keywords
Software Defined Network Exchanges (SDXs), Programmable networking and distributed clouds, Software Defined Networking, multi-domain networking, multi-service networking, multi-layer networking, network prototypes, GENI, iGENI, experimental network testbeds, programmable clouds

1. INTRODUCTION
Software-Defined-Networking (SDN) implementations have been proliferating among many types of networks, particularly within data centers and across private networks interconnecting data centers. SDNs have been particularly important for traffic management related to cloud based services. There are many advantages to SDN, including more options for dynamic provisioning options, faster response to changing traffic conditions, enhanced load balancing, more granulated traffic engineering, better network resource utilization, and enhanced opportunities to implement many new types of services. These and many other benefits are driving the deployment of SDN in all types of networks, from those in the enterprises to those provisioned by large scale providers.

The success of SDNs implementations has given rise to considerations of creating Software Defined Networking exchanges (SDXs) in order to provide SDN capabilities across multiple domains. This issue directly relates to the basic architecture of SDN, which is oriented to single domains – not multiple domains. However, these multi-domain considerations also give rise to multiple other challenges, including those related to signaling, messaging, resource advertising, description and tracking, mechanisms for topology exchange, multi-service provisioning, and others.

Today, SDXs are primarily conceptual, although several SDX prototypes exist and production SDXs are being planned. This paper describes a number of design considerations related to SDXs, as well as basic capabilities, underlying technologies and protocols, early prototypes, and initial experiences with one of those prototypes.

For example, the International Center for Advanced Internet Research (iCAIR) and its research partners designed and implemented a prototype SDX at the StarLight International/National Communications Exchange Facility, a major exchange facility for world-wide international, national, and regional research and education networks, data intensive science networks, federal agency networks, and large scale national and international network research testbeds (approximately 15-20 at any given time). This prototype SDX is being used for research experimentation and demonstrations. This SDX has been designed to be an international and national multi-domain, multi-service facility enabling federated controllers to manage network resources across existing barriers to extending control planes across WANs. The architecture, services, and technologies of such SDXs offer many complex challenges, but they also provide major opportunities for innovative networking, especially for programmable, segmented ultra high capacity networking required by next generation distributed clouds.

These opportunities are being provided by the rapid development of virtualization techniques that are enabling the implementation of significantly higher levels of abstraction for network control and management functions at all layers and across all underlying technologies. These approaches are not only allowing network designers to create a much wider range of services and capabilities

978-0-9836283-9-2

than can be provided with traditional networks, but they are also providing for a) many more dynamic provisioning options, including real time provisioning b) faster implementation of new and enhanced services c) enabling applications, edge processes and even individuals to directly control core resources; e) substantially improved options for creating customizable networks e) enhanced operational efficiency and effectiveness and f) more options for traffic engineering, especially those that are highly specific to particular types of traffic flows.

These capabilities are especially important resources for distributed clouds, particularly those that are distributed across multiple domains. Essentially, SDN enables a more optimal dynamic networking and matching of communication service requirements and network resources.

Although the many benefits of SDN are well known, particularly within and among large data centers, a number of issues remain to be resolved. For example, as noted, current SDN architecture is single domain oriented, and the required capabilities for multi-domain multi-service SDN provisioning remain fairly challenging. Another issue is that with the increasing deployment of SDN in production networks, the need for SDN exchanges has been recognized (SDXs). Today, although no such production facility exists, several are being planned, and prototypes have been implemented.

2. Software Defined Networking (SDN)

Before describing an SDX, it may be useful to provide an overview of SDN, a technique for programmable networking that, for the majority of implementations, is based on the OpenFlow. [1] Many new architectural techniques leverage capabilities for virtualization to allow for higher levels of abstraction for physical resources, configurations, and functions, e.g., the forwarding path. The SDN/OpenFlow architecture allows for an enhanced abstraction layer for network control, providing for a central overview of all major network components. Because the architecture separates the data plane from the control plane, provisioning, especially dynamic provisioning, is much more flexible. A controller is linked by a secure channel to an OpenFlow device, making it possible to externally manage network functions. For example, information on data flow behavior from flow tables in OpenFlow switches can be sent through the secure channel to the controller. When the controller receives the information, it can be analyzed, and, depending on the results of that analysis, traffic patterns can be directly and dynamically changed.

Until very recently, the SDN/OpenFlow approach has been used primarily within single domains, for example, within data centers, driven primarily by distributed cloud applications, and across data centers, supporting distributed cloud services. However, the success of the SDN/OpenFlow model has led to the consideration of extending this model across domains, especially to supports services based on distributed clouds. Consequently, several network research communities are beginning to investigate the concept of a Software Defined Networking Exchange (SDX).

3. SDX Services

As noted, SDN capabilities have been implemented successfully in single domains within data centers and between data centers. Therefore, a major design objective of a Software-Defined Network Exchange (SDX) is to provide for extensions for these SDN capabilities among domains. Such extensions of capabilities are particularly important for services based on multi-domain clouds. These capabilities include dynamic services provisioning, optimal resource utilization, load balancing, traffic conditioning, fault response, etc. Currently, these capabilities cannot extend across domains because communications exchange facilities such facilities are not enabled with SDN implementations. In order to provide for support of these capabilities across domains, a number of architectural issues must be addressed.

The foundation resources of an SDX consist of flexible infrastructure resources that are virtualized through various abstraction techniques, especially SDN/OpenFlow, enabling them to be highly programmable. Essentially, this approach creates a virtualized programmable environment, in which resources can be advertise, discovered, claimed, assembled, utilized and then discarded after use. Also, this environment enables specific application and services requirements can be precisely matched to network resources, while still using shared vs dedicated infrastructure. Beyond these considerations, a number of issues need to be addressed, including what resources can be found and claimed, what processes are used to advertise and signal for resources, what actual services can be dynamically provisioned, and how can controllers be federated to accomplish these processes.

Essentially, all exchanges provide mechanisms to allow for data traffic to flow among networks under various policy parameters. Today, the majority of data exchanges are L3 peering sites that depend on BGP to interconnect with other domains. The inherent limitations of BGP prevent more complex interdomain capabilities. Also, BGP does not address the need to support other service layers, e.g., L0, L1, L2, and hybrid services consisting of blended layers across domains. The potential for SDXs to support all service layers, including hybrid services, is particularly important for supporting distributed clouds. It is useful to note that to date almost all traditional L2 and OpenFlow services have been implemented as mutually exclusive options. More recently, network research communities have been designing and implementing techniques that enable both to not only co-exist but also to be integrated into a common service. Therefore, at an SDX, using these techniques, it would be possible to use traditional L2 services, or only OpenFlow services, and/or integrated L2/OF services.

4. STARLIGHT SOFTWARE DEFINED NETWORKING EXCHANGE (STARLIGHT SDXs)

The International Center for Advanced Internet Research (iCAIR) at Northwestern University and its research partners designed and implemented a prototype Software Defined Networking Exchange (SDX) at the StarLight International/National Communications Exchange Facility. The StarLight Facility, which has over 4 Tbps of capacity, in part, based on almost 30 100 Gbps paths, supports many different types of services for wide area private networks, including processes for implementing customizable services and core capabilities, for example, close integrations with specialized instruments, data intensive compute clusters, and high-performance data storage. [2] From its inception, StarLight was designed to enable traffic exchange at all layers, not simply at L3. To accomplish this goal, the facility has implemented techniques for abstracting control planes to allow for more programmability of services and capabilities, and for heterogeneous use of data planes, including pioneering techniques for dynamic lightpath provisioning. These techniques allow the high capacity channels at the facility, including multiple 100 Gbps, 40 Gbps, and many 10 Gbps optical lightpaths, to be partitioned and sub-partitioned across different domains, and each subset can be provided with individualized management and engineering functions.

StarLight is a particularly useful site for a prototype SDX also because it is a core resource for a large scale distributed global facility - the Global Lambda Integrated Facility (GLIF). (Ref Fig. 1) [3] The GLIF is not a network but more of a distributed facility within which it is possible to create customized networks. The GLIF interconnects a number of GLIF Open Lambda Exchanges (GOLEs) in North and South America, Asia, Europe and other parts of the world. StarLight is one of the GLIF GOLEs in North America. The GLIF is based on world-wide optical fiber implementations that support lightpaths that can be individual addressed, managed and controlled to create custom production networks and experimental testbed networks.

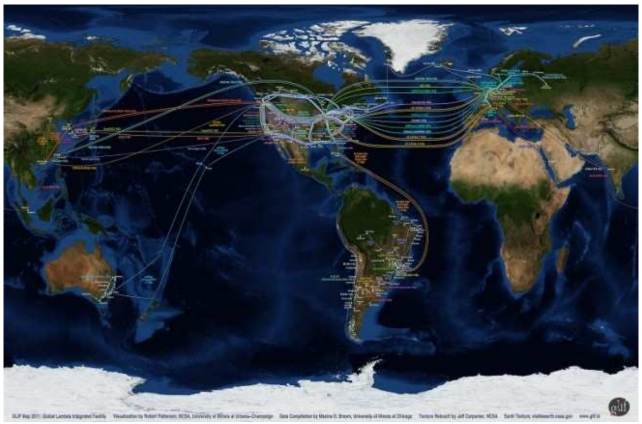

Figure 1: Global Lambda Integrated Facility (GLIF).

The GLIF community, in partnership with the StarLight consortium, other R&D communities, and the Open Grid Forum (OGF), a standards organization, have been developing an architecture for an interface (the Network Services Interface Connection Service or NSI-CS v 2.0 protocol) that can interoperate with the many individual control frameworks that have been implemented at the GLIF GOLEs in order to provide a common interface for multiple control frameworks for the dynamic provisioning of connection services. [4] Currently, several other initiatives are leading to several OGF NSI related documents that are being developed or are planned. These documents include those describing the NSI Framework, a standard, an NSI Network Service Agent Description Document, a standard, an informational NSI AAI (that is, an overview of Authentication, Authorization, Identification), a best practices description, an NSI Topology document, a standard, and an NSI Document Distribution Service (describing information distribution related to the NSI service), another standard.

These documents are noted here because they provide a useful context for many of the issues related to designing, implementing, and operating an SDX. To some degree, the NSI component developments anticipated the development of SDXs. It is also notable that the NSI-CS has been placed or is being placed in production within a number of international R&E networks.

The NSI Connection Service (CS) protocol provides for the reservation, creation, management, and removal of network connections, including across domains. Because the protocol assumes shared resources (e.g., circuits) across multiple domains, the NSI Connection Service incorporates capabilities for authentication and authorization.

In general practice today, service attributes are closely bound to the control and data planes that support them. NSI provides an abstract concept of a connection, which can enable connection services to transverse many domains. The NSI protocol includes a messaging specification that defines a schema that can be used to indicate service constraints that can be matched with local resources. A pathfinder process determines the proposed and then implemented path that will meet the service constraints. Consequently, one service plane protocol set of processes can provide for connections across multiple domains. The NSI protocol is essentially middleware between software requestor agents and software provider agents, using processes for federated resource discovery (Discovery Service) and topology sharing (Topology Service).

A more recent initiative is a project that has been established to integrate NSI with OpenFlow/SDN capabilities and to enable implementation of those capabilities at selected exchange points within the GLIF.

5. GLOBAL ENVIRONMENT FOR NETWORK INNOVATIONS (GENI)

Another reference architecture for the StarLight SDX is the National Science Foundation's Global Environment for Network Innovations (GENI) program, which has developed a nationwide distributed environment for experimental network research comprised of programmable networking resources closely integrated with programmable cloud resources. [5]

For example, the StarLight SDX incorporates many virtualization techniques used by GENI, including, Software as a Service (SaaS), Platform as a Service (PaaS), Infrastructure as a Service (IaaS), etc. As one example of using GENI as a reference architecture, the StarLight SDX uses the Flowvisor OpenFlow Aggregate Manager (FOAM), a mechanism with which it is possible to interact with other GENI domains as a resource allocation interface. [6] Also, the StarLight SDX uses the Floodlight and other Open SDN controllers, which are used in the GENI environment.

6. INTERNATIONAL GLOBAL ENVIRONMENT FOR NETWORK INNOVATIONS (iGENI) ANS ADVANCED INTERNATIONAL OPENFLOW TESTBED

Another reference architecture, and, to some degree, a prototype for the StarLight SDX builds on prior efforts that designed and implemented a large scale international experimental network research testbed based on OpenFlow/SDN, which incorporates also a distributed International GENI (iGENI) testbed. [7] For over four years, this world-wide testbed has been used by a global community of network research organizations to investigate various research topics by using the testbed as a highly partitionable assemblage of resources. (Ref: Figure 2)

Figure 2: International Advanced Networking Research Facility Based On OpenFlow/SDN

The international research community that has developed this OpenFlow based testbed has undertaken multiple experiments using it and has staged demonstrations at local, national and international forums, including the international SC, GLIF and GENI conferences. Topics have included many types of L2 functions including a POX based VLAN translation service, a NOX based multi-domain LLDP (Link Layer Discovery Protocol service), a NOX based OAM Continuity Check Message service (CCM), a Multipath TCP (MPTCP) integrated with Floodlight, prototype slice exchanges, and many other capabilities. [8, 9, 10, 11, 12, 13] (Ref: Figure 3).

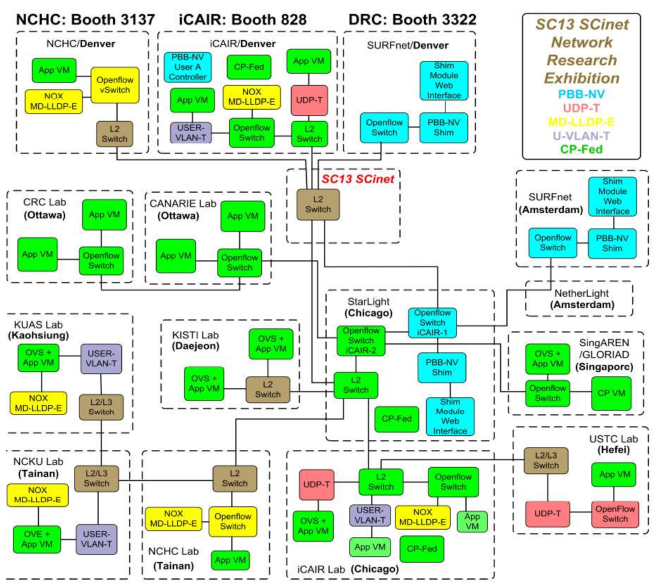

Figure 3: Software Defined Networking Exchanges (SDXs) Design Considerations

The overall architecture, implementation design, services, capabilities, and foundation technologies of SDXs are still being discussed and investigated. Today the majority of planned SDXs are oriented toward L3 services. [15] The SDX implementation at StarLight is a prototype that is oriented to L2 services. This SDX is intended to provide a means for migrating its services and capabilities to an ever larger set of virtual overlay networks – and even wider environments, including compute resources, clouds, storage devices, instruments, sensors, etc., basically, Software Defined Infrastructure (SDI). The SDX design incorporates high levels of programmability, open APIs, shared resources across multiple domains, dynamic provisioning, resource discovery, quick resource integration and configuration, and granulated control of resources. The StarLight SDX is a virtual, highly programmable facility, based on multiple addressable supporting resources (Ref: Fig 3). Abstractly, it can be considered a large scale virtual switch, which has resource elements that can be sub-partitioned and allocated to different domains.

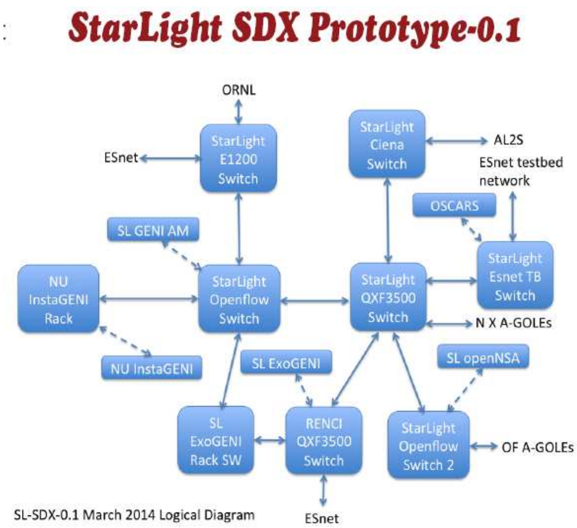

Figure 3: SDX Prototype

6.1 Path Controller

Inherent to the StarLight SDX is an SDN/OF/L2 integrated path controller, which functions as a link controller, and, thereby, allows other federated controllers to provide for resource management, as well as direct management by edge processes, such as applications, system processes, and even individuals directly managing network resources. As a result, this architecture enables such edge processes to optimize through signaling any type of service or capability, including resource discovery and integration, topology discovery, path selection and optimization, dynamic network resource provisioning, explicit data flow provisioning directly related to requirments, static resource provisioning, resource monitoring operations, fault detection and response, etc.

The StarLight SDX blends two architecture that have previously been provisioned separately, OpenFlow and traditional L2. L2 services can be directly integrated with aggregate managers and controllers in part because the SDN/OpenFlow architecture provides for separate control and data planes. This integrated design extends the benefits of the SDN/OpenFlow protocols to L2 networks, and it enables L2 benefits to be combined with SDN/OpenFlow based networks. In this architecture, using an OpenFlow controller provides a number of advantages, including enhanced interoperability among domains.

However, the overall design of the StarLight SDX provides additional capabilities – beyond those supported by standard OpenFlow implementations, including L2 functions. The design provides for a highly programmable platform comprised of many resources that can be controlled at a granulated level. Using this approach, edge processes can select from a number of potential services and options to more precisely match service and application requirements and network resources.

6.2 StarLight SDX Experiments and Demonstrations

Because the StarLight SDX has been designed as an international multi-domain, multi-service facility, it can allow federated controllers to directly manipulate network resources through control planes extended across WANs. Consequently, the StarLight prototype SDX is being used for many types of empirical research experimentation, for demonstrations, and for proof-of-concept showcases at various national and international workshops and conferences.

For example, it has been used to demonstrate the value of mechanisms for multi-domain automatic network topology discovery (MDANTD). [14] It has been used to demonstrate the advantages of content distribution networks that do not rely on IP addressing. It has been used to experiment with new types of OpenFlow controllers and controller functions. It has also been used to demonstrate large scale file transfers using multipathing with MPTCP [13], to showcase high performance transfers from specialized instrumentation, and capabilities for quickly provisioning customized specialized networks.

Recently, the StarLight SDX was used to support a demonstration staged at the GENI Engineering Conference in Atlanta (GEC 19), to showcase the advantages of a new type of severe weather prediction application. This application is being developed by the University of Massachusetts at Amherst. Current Doppler radar scans upward and cannot detect patterns over the curvature of the earth. The new application is based on small form factor Doppler radar that sites more instruments in more places to gather much more data to allow for highly granulated views of weather patterns. However, because this approach gathers so much data, it cannot be computed or stored locally. The data has to be transported in real time over L2 paths on optical fiber networks to remote cloud based compute sites where analytics take place and the results then are sent back to first responders. Essentially, the application requires a large scale virtual private network designed and implemented dynamically. For the GEC 19 demonstrations, multiple paths were dynamically instantiated across five separate domains, StarLight, the Southern Crossroads exchange facility (SOX), which is supporting another prototype SDX, and three separate optical fiber based networks each within its own domain between Chicago and Atlanta. (Ref Fig 4)

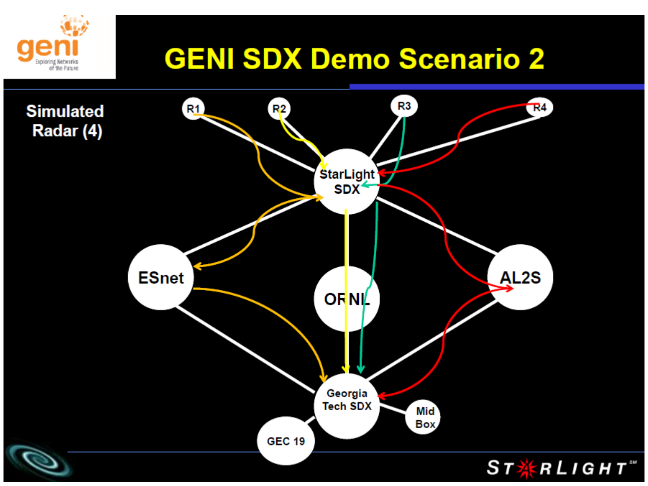

Figure 4: GEC 19 SDX Demonstration

Another set of demonstrations of international SDX dynamic provisioning was staged at the TERENA Networking Conference in Dublin, Ireland in May 2014 (Ref Fig. 5).

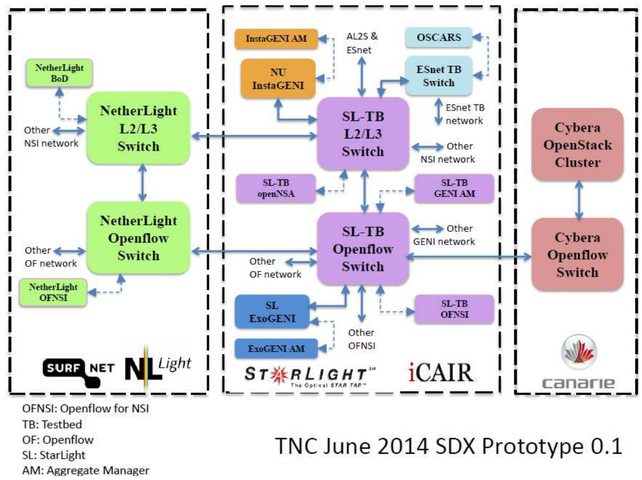

Figure 5: International SDX Demonstration at TNC 20014

7. SUMMARY

Over the last several years, Software Defined Networks (SDNs), almost all based on OpenFlow, have been deployed in single domain networks, especially driven by the needs of services supported by distributed clouds. The success of SDNs for many services and applications has led to the need to extend SDNs across multiple domains, to interconnect the growing number of SDN islands. Therefore, Software-Defined Network Exchanges (SDXs), are being designed and implemented in both prototype and in production. There are many alternative concepts being discussed and experimented with related to SDXs, which are still evolving. Initial SDXs are currently being used to conduct experiments, to demonstrate the potentials of SDXs, and they are also being developed for production.

There are many next steps that will be taken to advance SDX services and technologies. For example, currently, SDXs are based on OpenFlow architecture and protocols; however most networks are not. Consequently, an architectural approach is being development, in part by iCAIR, to enable integration of OpenFlow and non-OpenFlow networks. Another issue is the degree to which OpenFlow is an absolute requirement for highly programmable networks. Some current research investigations are exploring options as alternatives to OpenFlow. Another issue interoperability among different technical manufacturers' implementations of OpenFlow. To date, the current versions are interoperable, however as enhancements become more sophisticated, this may not be the case in the future.

One particularly interesting research area relates to using the granulated programmability of SDXs to support WAN service parameters for specific distributed applications, for example, by providing capabilities at SDXs to support "Application Exchanges" (e.g., multi-domain application specific peerings) using SDX services. For example, there have been recent discussions on this topic related to ultra definition digital media, 4k, 8k, ultra resolution 3D, all of which require high capacity individual streams. Another suggestion is providing ultra high segmented security trust-based channels through SDxs. Another is providing exchange based real-time transcoding for content distribution networks. Specific types of traffic and data can be highly segmented at a granulated level even across WANs [16] Because SDX are based on highly programmable resources through significant levels of virtualization, the potential for SDX based services is almost unlimited. In fact, SDXs can also be highly virtualized and distributed.

8. ACKNOWLEDGMENTS

For their first rate support for the initiatives described here, the authors would like to thank the US National Science Foundation, in particular, the Computer and Information Science and Engineering (CISE) Directorate, and its Global Environment for Network Innovations (GENI) program, including the International GENI (iGENI) program and the GENI Program Office the Office of Cyberinfrastructure and its International Research Network Connections (IRNC) program, the National Center for High-Performance Computing (NCHC), Taiwan, the National Kaohsiung University of Applied Science (KUAS), the National Cheng Kung University (NCKU) KUAS/NCKU, the National Taiwan University of Science and Technology (NTUST), the National Central University (NCU), SURFSara, SURFnet, the GigaPort program of the Dutch Government, the Geant3 project of the European Commission. Research at the Communications Research Centre of Canada. Canada's Advanced Research and Innovation Network (CANARIE), Cybera, the US Department of Energy's Office of Science, and members of the iCAIR research partnership consortiums, including members of the Global Lambda Integrated Facility (GLIF), the StarLight consortium, the Metropolitan Research and Education Network, the Energy Science Network (ESnet), GENI, International GENI, and the Open Cloud Consortium.

9. REFERENCES

[1] N. McKeown, et al., OpenFlow: Enabling Innovation in Campus Networks, ACM SIGCOMM Computer Communication Review, 2008, 2, 69-74.

[2] J. Mambretti, T. DeFanti, M. Brown, StarLight: Next-Generation Communication Services, Exchanges, and Global Facilities, Advances in Computers. 01/2010; 80:191-207.

[3] Global Lambda Integrated Facility (GLIF) www.glif.is

[4] G. Roberts, T. Kudoh, I. Monga, J. Sobieski, J. MacAuley, C. Guok, NSI Connection Service V2.0, Open Grid Forum, Group Working Draft (GWD), Candidate Recommendation Proposed (R-P), 2013.

[5] www.geni.net

[6] R. Sherwood, et al., "FlowVisor: A Network Virtualization Layer," OpenFlow Technical Report TR-2009-1.

[7] J. Mambretti, J. Chen, F. Yeh,International Network Research Testbed Facilities Based on OpenFlow: Architecture, Services, Technologies, and Distributed Infrastructure 18th IEEE International Conference on Networks (ICON), 2012, pp. 234 – 242

[8] M. Luo, S. Lin, J. Chen, "From Monolithic Systems to a Federated E-Learning Cloud System", IEEE International Conference on Cloud Engineering, March 25-28, San Francisco, California, USA, 2013.

[9] M. Luo, J. Chen, J. Mambretti, S. Lin, P. Tsai, F. Yeh, and C. Yang, "Network Virtualization Implementation over Global Research Production Networks", Journal of Internet Technology, Vol. 14, No.7, pp. 1061-1072, 2013.

[10] J. Mambretti, J. Chen, F. Yeh, T.-L. Liu, M.-Y. Luo, C.-S. Yang, R. v. d. Pol, A. Barczyk, F. Dijkstra, and G. v. Malensteinz, "OpenFlow Services for Science: An International Experimental Research Network Demonstrating Multi-Domain Automatic Network Topology Discovery, Direct Dynamic Path Provisioning Using Edge Signaling and Control, Integration With Multipathing Using MPTCP," in SCinet Research Sandbox at International Conference for High Performance Computing, Networking, Storage, and Analysis, White Paper and Presentation, SC12, November, 2012.

[11] M. Luo, J. Chen, "Towards Network Virtualization Management for Federated Cloud Systems", IEEE 6th International Conference on Cloud Computing, June 27-July 2, 2013, Santa Clara, CA, USA.

[12] M. Luo and J. Chen, "Software Defined Networking Across Distributed Datacenters over Cloud," 5th IEEE International Conference on Cloud Computing Technology and Science (IEEE CloudCom), Bristol, UK, December 2-5, 2013.

[13] R. van der Pol, Sander Boele, Freek Dijkstra, Artur Barczyky, Gerben van Malensteinz, Jim Hao Chen, and Joe Mambretti, Multipathing with MPTCP and OpenFlow, forthcoming, Proceedings SC12, 2012

[14] W-Y. Huang, J-W. Hu, S-C. Lin, T-L. Liu, P-W. Tsai, C-S. Yang, F. Yeh, J. Hao Chen, J. Mambretti, "Design and Implementation of An Automatic Network Topology Discovery System for the Future Internet Across Different Domains," Proceedings of IEEE 26th International Conference on Advanced Information Networking and Applications Workshops (AINAW'12), Singapore, March 2012.

[15] A. Guptay, L. Vanbever, M. Shahbazy, S. P. Donovany, B. Schlinkerz, N. Feamstery, J. Rexford, S. Shenker□, R. Clarky, E. Katz-Bassettz, SDX: A Software Defined Internet Exchange, accepted, ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking (HotSDN 2014), Co-located with ACM SIGCOMM'14, August 22, 2014, Chicago, IL.

[16] P-W. Tsai, P-W. Cheng, C.-S. Yang, M.-Y. Luo, J. Chen, Supporting Extensions of VLAN-tagged Traffic Across OpenFlow Networks, Research and Educational Experiment Workshop (GREE), 2013 Second GREE; 01/2013