OPERATIONALIZATION OF SOFTWARE-DEFINED NETWORKS (SDN) PROGRAM REVIEW

Dec 16-17, 2013

Conducted in Arlington, Virginia by

National Science Foundation

Department of Energy

National Coordination Office

U.S. DEPARTMENT OF ENERGY

NITRD

NSF

DISCLAIMER[1]

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

[1] Cover Photo from iStockPhoto

Editors

Inder	Monga	ESnet **(Chair)**
Grant	Miller	NITRD (Organizer)
Roy	Campbell	University of Illinois
Chip	Elliott	BBN
Ron	Hutchins	Georgia Institute of Technology

1 Executive Summary

U.S. federal investments in networking research and technologies deployment have fostered and accelerated the development of the Internet from its inception. It is now an essential infrastructure for the United States and the world. However, it is clear that we must make this infrastructure more flexible, resilient, secure, reliable, and ubiquitous to keep pace with society's needs and our aspirations for the 21st century.

Recently, U.S. agencies have been investing in networking innovation that will lead to the next generation of communication and cloud technologies to provide improved capacity, tools, service, and equipment needed for applications to be supported by the future Internet. New application requirements like those of large workloads moving between data centers, massive number of devices participating in machine to machine communication, and high-bandwidth applications on mobile devices are forcing change in network architecture to be more responsive rather than statically provisioned. One of these future network technologies, Software-Defined Networking (SDN), has the potential and momentum to create new engines of innovation and transform the entire Internet ecosystem. In December 2013, an invited review on operationalizing SDN was conducted at the National Science Foundation (NSF) with representatives from the academic, federal, and commercial communities. Workshop sponsorship was provided by the Computer and Information Science and Engineering Directorate (CISE) of the NSF and the Advanced Scientific Computing Research (ASCR) Program of the Department of Energy Office of Science, with support from the National Coordination Office for Networking and Information Technology Research and Development (NCO for NITRD).

Software Defined Networking (SDN) is an approach to computer networking architecture that radically decouples the system that decides where traffic is sent (the *control plane*) from the underlying systems that actually forward traffic to these selected destinations (the *data plane*). Through the logical separation of the network control and data planes, SDN technologies are enabling the creation of a new form of distributed infrastructure that can support advanced applications in the scientific, research and commercial world.

Many organizations in science, academia, and industry are experimenting with SDN technology, and working to quantify the benefits of SDN when it is applied or deployed in different contexts. Benefits range from lower capital and operational expenses to creation of new innovative new applications leveraging network programmability. For SDN to extend its impact to Internet scale there is now a critical need and an opportunity to extend SDN technology – both within a single domain and transparently across multiple domains like the Internet of today – with the intent to support novel SDN-based applications.

The SDN review brought together key individuals from the networking community to identify:

- Requirements, timing, and approach needed over the short term to develop, deploy and operate a prototype multi-domain SDN network supporting Internet-scale deployment of novel applications;
- Technological and operational gaps in need of research and development over the longer term to increase the capability of SDN networks, and better integrate them with existing public Internet and emerging cloud technologies;
- Security gaps and opportunities created by this approach, including opportunities to build in additional levels of security and resilience.

Key Findings

Many focused but fragmented efforts are underway including national R&E networks (ESnet, Internet2), advanced regional R&E networks, GENI and US Ignite, academic campuses, and the world's first SDN exchanges in the United States and New Zealand. There is also extensive industry activity in this space, ranging from active participation in large standards organizations to collaborative open-source efforts and deployments in data centers and wide-area networks to Network Functions Virtualization in telecom service providers.

Using such SDN technology, we can now envision (and in practical terms, create) scientific "instruments on demand" or application-specific "infrastructure on demand" across multiple networks (multi-domain), on a worldwide scale. Nearly all participants expressed an interest in developing end-to-end services spanning multiple SDN networks. Also, prototyping was seen as a viable means to provide the experience needed to develop a pragmatic, scalable, and sustainable approach to multi-domain SDNs. There was consistent acknowledgement that the time is right for deploying prototype operational, multi-domain SDNs

Despite its clear potential for high payoff and enthusiastic vendor adoption, the requisite SDN hardware and software technology is still immature, with key aspects still in the research phase. Multi-domain software-defined networks are not yet operationally deployed and it likely will take several iterations of design and experimentation before we have systems that work well in practical, operational terms. The currently available SDN hardware bookends the possible capability spectrum with one end focused on the data-center oriented, white-box switches and the other end on highly capable, expensive, programmable routers. There is ample opportunity to significantly improve the hardware and software offerings in the marketplace through active experimentation and prototype deployment.

The security aspects of SDN-based infrastructure need additional research, development, and experimentation to ensure successful deployment of SDN technology at multiple levels of abstraction.

The security-related research topics include requirements, trust models, and attack scenarios. In addition, more practical concerns of resilient design, authentication, authorization, asynchronous operation, maintenance, and software engineering are in need of research and development. This is true of both single-domain and multi-domain SDNs.

In addition, adoption of SDN is expected to open up possibilities to exploit the new abstractions and programmability in security policies and security applications. Areas of development include secure updates for applications, flexible intrusion detection, and flexible reaction and provisioning.

Key Recommendations

The SDN workshop participants discussed solutions to the gaps and opportunities for impact and shortlisted a few actionable recommendations to accomplish the goals. At a high level, key representative recommendations include:

1. **The United States Government (USG) agencies should sponsor efforts to research, design, deploy, and operate prototype multi-domain** software-defined networks **as soon as possible**, where SDN is understood as enabling the entire distributed infrastructure needed for next-generation commercial and/or scientific applications. The focus of these efforts would not only be the enabling of new end-to-end applications, but also the development of necessary operational tools needed to manage and operate software-defined networks in production.

   The SDN ecosystem (SDN-based Internet or S-Net) will need close ties to the commercial sector, and active participation from researchers, applications and instrument engineers, and network and software engineers so the participants would learn quickly from that experience and have the ability to innovate and implement the operational insights.
2. **An initial focus is needed on operational SDN deployments including Software-Defined Exchanges (SDXs)** to enable interoperability and use of these new approaches with the current Internet infrastructure. Initially, these efforts should be focused on defining the architectures/implementations that will support operational multi-domain SDNs; experimenting with these designs and iterating as necessary; encouraging the development and deployment of next-generation instruments and applications made possible by this infrastructure; growing the community of multi-SDN aware engineers, researchers, and students; and preparing for transition to an operational phase.
3. **Since cyber-security is of the highest importance for deployed, multi-domain, multi-layer software-defined networks**, a vigorous and sustained research program should investigate the security implications. This research will benefit

from the close interactions of security researchers with the engineers and operators of the deployed multi-domain, multi-layer SDN prototypes, and with engineers of the applications or instruments that are supported by those networks.

4. **Enabling the integration of the network with the other elements of the software-defined infrastructure, namely, compute, storage and sensors**. In an environment where networks are responsive to application needs through their programmability, effort should be taken to open the network 'black box' i.e. to expose and integrate the informational and programmable elements into the larger infrastructure ecosystem enabling easy orchestration of resources by applications.
5. **Investing in tools and procedures for managing operational SDN networks** was a strong recommendation. This lack of functionality is typical of new technologies and participants felt that the only way to bridge the gap was to invest in building tools and leveraging community best practices and knowledge that will enable viable operational models for accelerating deployment of this technology.
6. **A coordinating effort is needed for capturing and sharing gained practical knowledge to educate the community** in building and operating SDN networks as an important element for providing longer-term impacts. These educational efforts would be less oriented towards the technology and research understanding of SDN but be oriented towards disseminating the community best practices and knowledge gained from operating the prototype infrastructure.

Conclusions

Based on extensive discussions and key presentations, the participants identified the review as very timely, and strongly endorsed the need to help SDN evolve from its current state of tremendous promise to an operationally sound technology by building secure, prototype SDN network deployments. In addition, they identified end-to-end service manageability and security as key research and development areas that will be critical to enabling wider deployment. Building pragmatic approaches to seamlessly interoperate SDN domains with the existing Internet through software-defined exchanges was seen as an important first step towards realizing the goals and conclusions of the program review. A draft roadmap on how the SDXs can be deployed is instantiated below.

Draft Roadmap for Software-Defined Exchanges

The participants suggested a draft roadmap to indicate how these recommendations can be instantiated for experimentation or deployment.

Year	Goals, participants, activities
1	**Goals**: Stand up earliest versions of SDXs, develop and demonstrate early instruments and next-generation apps **Participants**: US Government agencies and programs (NSF, DOE, GENI, US Ignite, ESnet, DREN), academic institutions, regional and national networks (like Internet2) and collaboration with commercial exchanges and vendors. **Activities**: Identify participants (both for SDXs and early apps/instruments), create a forum for community participation, establish outreach and technical interchange activities, engage with US companies and US government agencies, begin security discussions, buy equipment and create software as needed to stand up early SDXs, build early prototypes of tools, port apps / instruments to work across multi-domain SDN, and give demo(s). Create security assessments, risk models, attack scenarios, auditing, and component vulnerability analysis.
2-3	**Goals**: Refine understanding of SDX technology including security and multi-domain aspects, increase heterogeneity, add participants, and begin standardization/interoperability **Participants**: Security researchers and professionals, additional US agencies (e.g., NIH, NASA, NOAA, ...), and more US operators & vendors, researchers **Activities**: Document the multiple competing SDN technical approaches and discuss them, start security analyses and red teams, identify tool chain requirements, add many new U.S. agencies and companies as participants, add several additional apps/instruments, ramp up researcher funding, begin design of curricula, discuss operational needs and potential forums for operator interactions, possibly add more SDXs, and perform early demos of interoperability between competing SDX approaches. Perform security analysis of Multi-Domain, Multi-Tenancy, and Denial of Service concerns. Investigate security policies and mechanisms for defense in depth and novel SDN approaches.
4	**Goals**: Transition at least 2-3 SDXs to full operational capability **Participants**: U.S. vendors, U.S. network operators (R&E, commercial), wide community of application/instrument developers, educators & students

	Activities: Agree on and publish basic "standards" for interoperability, harden the most useful tools in the tool chains, document security findings and recommendations, transition to SDN operations, conduct classes and training sessions using this new technology, and continue to enlarge both the infrastructure and suite of apps/instruments. On an ongoing basis, perform operational monitoring, alarms, audit controls, and reactive defense, building upon lessons learned to date with GENI.

2 Workshop Background, Goals, Structure and Timeline

A SDN approach to building networks holds promise for making the networks responsive to and meeting the varied needs of numerous applications. In addition, SDN programmability enables customization of the network suited to a particular application. This powerful paradigm and technology is currently being tried out in many domains by many organizations, with each domain investigating its own flavor. Each organization has created customized and application-specific software to adapt SDN to their needs. For SDN to extend its impact by becoming production ready and deployed at Internet-scale, now is a critical time for stakeholders to consolidate efforts and work collaboratively to ensure SDN implementations are capable of not only working within a single domain but also transparently across multiple domains.

This SDN review was held to identify the near-term requirements, processes, and players to develop, deploy, operate, and manage a prototype SDN network and, over the longer term, provide the research and development needed to extend SDN capabilities for end-users.

Workshop Background

Within this context, the Office of Science and Technology Policy (OSTP) directed the federal agencies participating in the Networking and Information Technology Research and Development (NITRD) Subcommittee Large Scale Networking (LSN) Coordinating Group to plan and hold a program review with participation by representatives from federal agencies, the commercial sector, researchers, and other networking and distributed systems research community participants to explore and report on the need for a prototype SDN network.

Topics of discussion included:

- Identifying current capabilities and resources that contribute to the development and operation of an operational SDN prototype network that interoperates seamlessly with the current public Internet;
- Identifying the research, resource, and collaboration needs for creating such a system;
- Identifying gaps in operational software tools necessary for running a production multi-domain SDN network;
- Identifying opportunities for SDN virtualized networks to interact with “clouds” of computation and storage; and
- Providing a workshop report to the NITRD Subcommittee and OSTP on recommendations for needed R&D, resources and collaboration for the prototype system.

To this end, the LSN agencies held an SDN prototype operational networking review on December 17–18, 2014, at the National Science Foundation with representatives from the academic, federal, and commercial communities. Funding was provided by the National Science Foundation Computer and Information Science and Engineering Directorate (CISE) and the Advanced Scientific Computational Research (ASCR), Department of Energy Office of Science, with support from the National Coordination Office for Networking and Information Technology Research and Development (NCO for NITRD).

Workshop Goals

The goals of the workshop were to bring together key individuals to:

1. Identify the requirements, timing, and responsibilities needed, over the short term, to develop, deploy, and operate a prototype multi-domain SDN network with:

- Transparency and interoperability among SDN domains,
- Transparency and interoperability with the public Internet,
- Acceptable levels of cybersecurity and robustness,
- Technology development for Layer-1, Layer-2, and Layer-3 operations,
- Inter-domain policy issues (control, identity management, information sharing, etc.),
- Advanced network capabilities demonstrations with SDN-like quality of service, efficient equipment usage, and energy use reductions,
- Participation of commercial equipment providers to facilitate technology transfer, and
- Development of new applications that leverage novel SDN capabilities.

2. Identify needed research and development, over the longer term, to increase the capability of SDN networks to support user applications, and to better integrate SDN technologies with the public Internet and emerging cloud technologies.

3. Provide a workshop report documenting recommendations for needed research and development, resources and collaboration to deploy and operate the prototype nationwide SDN network and to identify future SDN research needs.

Workshop Structure

The workshop organized plenary presentations to set the stage for common understanding and discussion in subsequent breakout sessions. The plenary presentations included:

- **Macro Trends, Complexity, and Current Status of SDN** by David Meyer
- **Future of SDN** by Jennifer Rexford
- **Building and Deploying SDN** by Lorenzo Vicisano
- **Innovation in Academia: Deployment, Operations, and Management of SDN** by Rob Vietzke
- **Security and SDN** by Roy Campbell

Additional lightning talks addressed Software Defined Internet Exchanges (SDXs), SDN research in production, CORONET, Network Visualization, and Multi-layer/ Multi-domain SDN. These talks provided the context for discussion in the workshop breakout groups. There were three breakout groups:

- **Users, Applications and Motivation** led by Chip Elliott and Ron Hutchins
- **Technology and Operational Gap Analysis** led by Inder Monga, Bill Snow, and Eric Boyd
- **Security and Policy** led by Roy Campbell

This report documents the findings and recommendations as developed in these breakout sessions.

3 Key Observations and Findings

The following section summarizes key observations and findings from the discussions within the breakout groups, and common points brought up across multiple breakout sessions.

3.1 General Findings

Out of the workshop sessions, there were many recurring themes, concerns, and ideas surrounding SDN. The main findings are highlighted below.

1. The time is right for building prototype, operational, multi-domain software-defined networks.

2. SDN should be thought of as encompassing the entire distributed infrastructure needed for science—i.e., a close integration of resources including compute, storage, and networks, building-in appropriate middleware, and security with energy efficiency considerations.

3. With SDN, we can now envision (and in practical terms, create) scientific "instruments on-demand" or application-specific "infrastructure on demand" on a multi-domain or worldwide scale.

4. Despite its clear potential for high payoff, the requisite SDN technology is still immature, with key aspects still in the research phase. Multi-domain/multi-layer software-defined networks do not yet exist and it will likely take several rounds of design and experimentation before we have systems that work well in practical, operational terms.

5. Because the operational implications of multi-domain/multi-layer SDNs are not yet clear, network engineers and operators will need to gain experience from running prototype multi-domain SDNs before they have sufficient knowledge and experience to reliably operate production multi-domain SDNs.

6. Also because the security implications of multi-domain/multi-layer SDNs are not yet understood, significant research efforts will be required to gain this knowledge. These security research areas will benefit from interactions with prototype multi-domain/multi-layer SDNs that will serve to motivate and focus the research.

7. Vigorous efforts will be needed to establish and grow the human community needed to successfully achieve multi-domain/multi-layer SDNs, including but not limited to: active creation and sustainment of such a community, technical

training for infrastructure engineers, college and university education, and strong industry-academic collaboration.

3.2 Users and Applications – Findings

The findings in this section arose from the Users, Applications, and Motivation break-out session. Please see Section 5.1 for a narrative description of the UAM break-out session.

1. To create new instruments and infrastructures enabled by SDNs, innovators of the future will build on the following foundations: (i) state-of-the-art, evolvable engineering systems and their tool chains, (ii) organizational cultures that understand and enable this new world, (iii) the people and processes that will create this new paradigm, and (iv) deeply multi-disciplinary approaches.

2. Productivity in such an environment will require the development of tool chains, middleware, and intermediate platforms that will support efficient application and instrument development and deployment. These software components will also help to build-in appropriate notions of security, service-tailoring, and energy efficiency.

3. Research and discovery enabled by such new infrastructure will include many types of networking/distributed systems research (e.g. virtualization, global orchestration, high availability, optimizations, machine learning and reasoning, customized services and networks, etc.), domain research (creating "world instruments" on demand in a range of domains, global resource optimization, etc.), and promising new areas of cross-cutting research (such as the interactions of SDN and wireless technology, policy, economics, etc.).

4. Applications enabled by such new infrastructure will include applications in healthcare, education, transportation, public safety, and advanced manufacturing enabled by deterministic quality of service. Parameters of interest to the application will include latency, jitter, flow rate, redundancy, reliability, security, cost, and availability of in-network computing and storage services.

5. Longer term, the group envisions "app stores" that include SDN-enabled software, which would open up the network infrastructure market to many actors, as has already occurred with smart phones. Such app stores could have

profound technological and economic implications, and form a fruitful area for research and experimentation, as well as economic development.

3.3 Technology and Operational Deployment – Findings

The findings in this section arose from the Technology and Operational Deployment (TOD) break-out session. Please see Section 5.2 for a narrative description of the TOD break-out session.

6. The models for interactions between applications and the network, including trust and security, need to be explored in order to increase the value of programmability and the creation of SDN-aware applications. A definition of the ecosystem for applications that includes the interaction between storage, compute, and network needs to be addressed.

   Capable, experimental platforms are needed that could be based on open source hardware and software to enable network operations and application engineers to experiment with the various SDN concepts so they can learn and adopt the new paradigm of programmability.

7. Design patterns and best practices for building SDNs have not been established because operational experience with these networks is lacking. The community should build interoperable solutions using existing standards to gain experience for establishing such rules of thumb for networking with this new paradigm.

8. Deployment of a secure and stable control plane is a critical concern in the deployment of an SDN network. Clear understanding and best practices for deploying and managing that control plane needs to exist in the community to facilitate adoption of SDN.

9. There is a large gap in the tools and understanding of how SDN/OpenFlow networks can be managed from a network operator perspective. Investigation into and investment in network monitoring, management and analysis tools is needed in order to increase the ease of deploying SDN networks.

10. Even though the concept of SDN controllers[2] for a Network Operating System has been established, the lack of standard APIs and a set of accepted functionality associated with a controller make the adoption of SDN harder because of the differentiation between various closed or open source controllers. The adoption of standard abstractions will encourage and sustain innovation both in the design and implementation of SDN controllers and within the network application domains as well.

11. Ultimately, a well understood migration path from the existing network models to SDN/OpenFlow-based networks is needed.

12. The current programmability of APIs is insufficiently abstracted to allow applications to interact effectively with the network—more effort needs to be put into building the right abstractions including monitoring and control to not only enable application programmability but also to allow network operators/engineers to manage that interaction.

13. Flexible description of application profiles is needed to handle the different kinds of applications since commercial application characteristics may vary significantly. Not all applications need to be network-aware or be able to program the network. These profiles need to be developed jointly by application and network engineers, and can be application specific.

14. By managing the underlying network, the physical network with multiple layers (like optical and Ethernet) across different physical media needs to be exposed to the higher logical layers that enable applications to build flexible overlay networks over the current physical infrastructure.

3.4 Security – Findings

The findings in this section arose from the Security break-out session. Please see Section 6.3 for a narrative description of the Security break-out session.

The architecture of a Software-Defined Network introduces both significant benefits and significant concerns for security. In addition, rapid pace of academic research,

[2] The term 'SDN Controller' is used in the document with a broad definition of set of software entities controlling the network. Other network functional entities, like Resource Brokers, could also be considered controllers or part of the control hierarchy in this context.

innovation, and open-source software development are introducing new approaches that require a matched pace of architectural understanding and review from a security perspective. The security implications of SDNs are poorly understood at present, and many important security questions arise even within the context of a single SDN - for instance, how can we assume the connections are free from tampering? In addition, in an SDN, we also have to consider the relationship between switches and the controller: How is trust established as the network, or parts of the network, are (re)booted. How should a switch trust all commands it receives from a controller? How should the controller trust information from the switches? How can a controller protect against impersonation of one switch by another or by a non-switch entity as a switch? How do controllers and switches need to authenticate and authorize one another?

15. Multi-domain SDN infrastructure must deal with the real-world complexities presented when a network comprises of distinct autonomous domains. From a security perspective, the key properties of such a system include: (a) the presence of a large number of distinct autonomous domains, and (b) an explicit lack of trust between all, "the entirety", of the autonomous domains.

16. Research should help define explicit mechanisms and interfaces by which application-level protocols and services may expose information about principals, suitable for authentication, authorization and policy enforcement without compromising the overall security of the network.

17. Because SDN controllers are highly likely to run not just one monolithic program but a collection of "apps" representing differing functionality, research is needed into the forms of SDN modularity that will enable us to safely compose these applications in predictable ways.

18. SDN security architecture and mechanisms must be designed for (eventual) scalability.

19. SDN security policies must support pairwise and explicit negotiation of security services – information visibility, mutual (pairwise) policy enforcement, while enabling and supporting security services for end-users that span the entire SDN infrastructure ("SDN-wide security objectives"). This is in direct opposition to a simpler model, in which all autonomous domains, each running a single local SDN in an inter-domain SDN eco-system, sign-up and agree to a common, universal set of security policies and mechanisms.

20. Allowing SDN-enabled security provisions to exploit the new abstractions and programmability of the networks introduces the possibilities for research to use these technologies in security policies and security applications. Such research might examine secure updates for applications, flexible intrusion detection, and flexible reaction and provisioning.

4 Recommendations

The following section captures the actionable recommendations in order to implement the goals of the workshop

1. **The United States Government (USG) agencies should sponsor efforts to research, design, stand up, and operate generic multi-domain SDNs as soon as possible**, where SDN is understood as encompassing the entire distributed infrastructure needed for next-generation commercial and/or scientific applications – i.e., closely integrated resources: compute, storage, and networks taking performance, reliability, security and energy economy into account.

   Creation of multiple SDNs connected with each other is needed to facilitate an environment where new approaches towards operational and management tools can be experimentally deployed, vetted and improved. This infrastructure and community engagement will be critical in moving this new technology from research into wider production. The SDN ecosystem (SDN-based Internet or S-Net) will need close ties to the commercial sector, and active participation from researchers, applications/instrument engineers, and network, and software engineers so the participants would not only learn quickly from that experience but also have the ability to innovate and implement the operational insights.

2. **The focus of initial SDN deployments should include Software-Defined Exchanges (SDXs)** to enable interoperability and use of these new approaches with the current Internet applications. Needs include:

   a) A small number (e.g. 3-5) of open Software Defined Exchanges (SDXs) within the United States to (a) interconnect peer SDNs, and (b) connect customers to these inter-connected SDNs. These SDXs should be co-designed in close collaboration with U.S. industry.

   b) Multiple parallel efforts to create, debug, and publish open source software that will make it easy to stand up end-to-end services between SDN islands, and to integrate this software with SDN networks, SDXs, and

next-generation instruments and applications. A key goal should be to create state-of-the-art, evolvable engineering systems and their tool chains, as well as the environment in which these tools are used and evolved.

c) An explicit effort to coordinate these activities and to establish and grow the human community needed to successfully achieve multi-domain SDNs, including, but not limited to: active creation and sustainment of such a community, technical training for infrastructure engineers, college/university education, and strong industry-academic collaboration. This applies equally to the SDX sites, the tool-chain (see Section 6.1 for a description of the tool chain), middleware, and platform communities, and the scientific instrument and next-generation applications communities.

d) Initial trials that begin as soon as possible, and include the existing SDN exchanges within the U.S., R&E networks like ESnet and Internet2, and those regional R&E networks and campuses that are prepared to participate. Multi-domain/multi-layer concepts and potential applications should be demonstrated via early prototypes, as soon as possible, to give practical insights into what will be needed, and to help give the networking/applications communities a concrete sense of what is being planned

For the first few years, these coordinated efforts should be focused on defining the architectures/implementations that will support operational multi-domain SDNs; experimenting with these designs and iterating as necessary; encouraging the development and deployment of next-generation instruments and applications made possible by this infrastructure; growing the community of multi-SDN engineers, researchers, and students; and preparing for transition to an operational phase. To encourage rapid progress in creating an active, open community:

- These efforts should actively engage key scientific instruments and next-generation applications as design and prototyping partners to ensure that the SDX design properly supports advanced instruments and applications. Each of the following areas should be represented by at least one such instrument/application: scientific instruments on demand, infrastructure on demand (e.g. near-term weather prediction), applications that provide bandwidth reservations and path optimization, and US Ignite applications.

- In general, these SDXs should be operated by organizations that are distinct from the SDN network operators, to ensure maximal inter-operable technical development and growth in the existing and future

SDN communities. However, regional networks may operate some SDXs.

- These SDXs should be co-designed with U.S. industry. Since companies may act competitively in this technology space, one approach would be to enlist one U.S. company or a small complementary set of companies per SDX, which would allow a number of competing ideas/designs to be tried in parallel on the multiple SDX prototypes.

- The first SDXs should include exchanges that can readily connect as peers (a) U.S.G SDN networks, (b) international SDN networks, and (c) existing R&E and US Ignite networks with SDN capability.

- It is desirable that international SDXs be included in the planning from the start, as this will give key insights into both the potential of, and technical and non-technical issues concerning, global-scale multi-domain SDNs.

3. Since **cyber-security is of the highest importance for deployed, multi-domain/multi-layer SDNs**, a vigorous and sustained research program should investigate the security implications of multi-domain/multi-layer SDNs. This research will benefit from close interactions of security researchers with the engineers and operators of the deployed multi-domain/multi-layer SDN prototypes, and with engineers of the applications /instruments that are supported by the multi-domain/multi-layer SDN.

4. **Enable the integration of the network with the other elements of the software-defined infrastructure, namely, compute, storage and sensors**. In an environment where networks are responsive to application needs through their programmability, effort should be taken to open the network "black box" i.e., to expose and integrate the informational and programmable elements into the larger infrastructure ecosystem enabling easy orchestration of resources by applications.

5. **Invest in tools and procedures for managing operational SDN networks** was a strong recommendation. This lack of functionality is typical of new technologies and it was felt that the only way to bridge the gap was to invest in building community best practices, knowledge and investing in tools that would enable viable operational models for accelerating deployment of this technology.

6. **Coordinate an effort for capturing and sharing gained practical knowledge to educate the community** in building and operating SDN networks as an important element for providing longer-term impacts. These educational efforts should be

less oriented towards the technology and research understanding of SDN and more oriented towards disseminating the community best practices and knowledge gained from operating the prototype infrastructure.

5 Perspectives from the review

The following section describes in detail some of the discussions within the workshop's breakout sessions.

5.1 Users and Applications

The Users, Applications, and Motivation (UAM) breakout session consisted of 20+ individuals representing a range of research domains, testbeds, network operators, applications, and equipment manufacturers. All had considerable practical familiarity with SDN. In the end, the group came to fairly widespread consensus on the topics presented below.

Where are we today?

There is tremendous industry activity in the SDN space, ranging from standards organizations to SDN/OpenFlow, OpenStack and other similar efforts in data centers and wide-area networks to Network Functions Virtualization in telecom service providers. However, most current SDN work is inwardly focused to help network providers/operators. Furthermore, it currently exists in stand-alone "islands" that cannot yet be interconnected to support multi-domain applications. That said, we are now starting to see interesting new trends towards SDXs where new technology at a single SDX can yield benefits for 10s or 100s of connecting networks. Early versions of SDXs are now being instantiated in Atlanta (SOX), Berkeley (ESnet), Chicago (Starlight), and New Zealand (REANNZ).

From an application viewpoint, the term "SDN" should be thought of as encompassing the entire distributed infrastructure needed for a scientific instrument / application – i.e., a close integration of resources: compute, storage, and networks, building-in appropriate middleware, security, and energy-efficiency considerations. Using such SDN technology, we can now envision (and in practical terms, create) scientific "instruments on demand" or application-specific "infrastructure on demand" across multiple networks (multi-domain), on a worldwide scale.

Given this broader SDN perspective, a variety of existing applications are ready to take advantage of SDN capabilities once they can be offered "end to end" across multiple SDN islands. Such applications range from bandwidth calendaring, scheduled data transfer, and advanced forms of network function virtualization, to CASA weather radar systems (on-demand clouds) and global resource optimization (e.g. Belle II collaboration).

Where could we be soon?

In the network-operator context, we can think about new SDN-enabled services, such as application-specific peering (e.g. for video), redirection to middle boxes, traffic offloading, and prevention of free riders. These new services are likely to reduce

expenses and improve efficiency for network operators, but by themselves will offer few if any new benefits to the applications riding atop such infrastructure.

From the perspective of scientific instruments or applications, however, we can begin to create quite interesting new kinds of applications/instruments once SDN capabilities are fully embraced, such as migration of virtual machines towards data, ephemeral scientific instruments, new approaches towards building highly resilient systems from cheap equipment, on-demand and highly responsive cyber-infrastructure, and so forth. The combination of SDN and wireless looks as if it could be particularly fruitful for applications, as it could enable such new fields as on-demand sensor networks.

Applications enabled by such new SDN infrastructure will include applications in healthcare, education, transportation, public safety, and advanced manufacturing enabled by deterministic quality of service. Parameters of interest to the application will include requirements for latency, jitter, flow rate, redundancy, reliability, security, cost, and availability of in-network computing and storage services. Research and discovery enabled by multi-domain SDNs will include many types of networking/distributed systems research (e.g. virtualization, global orchestration, high availability, optimizations, machine learning and reasoning, customized services and networks, etc.), domain research (creating "world instruments" on demand in a range of domains, global resource optimization, etc.), and promising new areas of cross-cutting research (such as the interactions of SDN and wireless technology, policy, economics, etc.).

Longer term, once we have semantic descriptions of applications/instruments and available SDN infrastructure, we can begin to perform machine reasoning and machine learning to create and continuously optimize new instruments and applications. Although this is still within the realm of research, its practical benefits would be very high, as it would permit automated creation of scientific instruments and highly individualized applications.

What additional tools and capabilities are needed?

Although we are already in a position to begin prototyping applications that take advantage of SDNs, many additional tools and capabilities need to be created, including:

- Techniques for setting up slices across multiple SDN domains
- Robust isolation of flows (performance, security) with guarantees and enforcement
- Safe forms of delegation
- Example "Hello world" applications that can be adapted as needed, with a range of default "zero knob" starter applications for various fields
- Tailored visibility of applications into cross-domain infrastructure, e.g., for debugging and performance improvements
- Easy-to-understand dashboards
- End-to-end debugging tools
- Simple core design/reference platforms with widely accepted APIs

- DevOps tools
- High-level languages for SDN description and manipulation
- Tools for manipulating semantic SDN descriptions with multiple levels of abstraction

Productivity in such an environment will require the development of tool chains, middleware, appropriate abstractions, and intermediate platforms that will support efficient application and instrument development and deployment. These software components will also help to build-in appropriate notions of security, service-tailoring, and energy efficiency.

Tools and tool chains for users (broadly defined) will be extremely important. We envision that some of these users will be actively engaged in creating and maintaining the requisite tool chains, most of which will probably be open source. Other users will simply use the tools. Key parts of the tool chain ecosystem will include:

- Basic debugging tools – need decomposition of responsibilities
- Embedded instrumentation
- Zero-knob slice instantiation
- Fault/delay injection
- Flow visualization
- Contention management
- Virtual perfSonar
- Slice management
- Open source whenever possible
- Strong O&M tools
- DevOps tools for all
- Transition tools

Longer term we envision SDN app. stores which will contain a very large number of third-party applications that can be easily instantiated by end users within the multi-domain SDN infrastructure. Although this is a fairly distant goal, it is worth working towards, given the major effects that app. stores have had to date for smart phones.

How do we get there?

To create new instruments and infrastructures enabled by SDNs, innovators of the future will build on the following foundations: (i) state-of-the-art, evolvable engineering systems and their tool chains, (ii) organizational cultures that understand and enable this new world, (iii) the people and processes that will create this new paradigm, and (iv) deeply multi-disciplinary approaches.

Since SDN technology is still in its infancy, many aspects of how applications and instruments take advantage of SDN are still unclear. The best way to proceed in such a situation is to create a prototype multi-domain SDN infrastructure as soon as possible

and begin experimentation with applications and instruments that take advantage of its capabilities. Along the way, we will start to build the new tool chains needed for such applications. This will require deep, prolonged interaction between the SDN infrastructure community (researchers and engineers) and several "early adopter" application and instrument communities. Very likely our understanding will evolve considerably during this process, and we should plan to revise both infrastructure and applications several times before optimal combinations of tools and capabilities are achieved.

Vigorous efforts will be needed to establish and grow the human community needed to successfully achieve multi-domain/multi-layer SDNs, including but not limited to active creation and sustainment of such a community, technical training for infrastructure engineers, college/university education, and strong industry-academic collaboration.

SDXs will provide an excellent focal point for creating and growing this new community. In our opinion, "private peering" between SDNs will provide few chances for enlarging the community of knowledgeable researchers, engineers, and scientific instrument builders; by contrast, the SDX approach makes it very easy to constantly grow the community by introducing new participants and bringing them up to speed.

5.2 Technology and Operational Deployment

The Technology and Operational Gaps (TOD) break-out session consisted of 40+ individuals representing a range of research domains, testbeds, network operators, and equipment manufacturers. The group raised the following issues on the topics presented below.

Current SDN Limitations

Design patterns for building SDNs have not been established because operational experience with these networks is lacking. The community should build interoperable solutions using existing standards and gain experience to establish such rules of networking with this new paradigm.

Even though the concept of SDN controllers and Network Operating Systems has been established, the lack of standard APIs and a set of accepted functionality associated with a controller make differentiation of various closed or open-source controllers and ultimately, adoption of SDN harder.

A well-understood migration path from existing paradigms and networks to SDN/OpenFlow-based networks is needed.

Understanding, representation and discovery of underlying network topology is extremely important to automation and programmability promised by SDN. Lack of standards, tools and understanding within the industry needs to be addressed, and awareness of this issue needs to be raised.

Operational Issues

There is a large gap in tools and understanding on how SDN/OpenFlow networks can be managed from a network operator perspective. Investigation and investment in network monitoring, management and analysis tools is needed in order to improve the ease of deploying SDN networks.

Deployment of a secure and stable control plane is of critical concern in deployment of an SDN network. Clear understanding and best practices for deploying and managing that control plane need to exist in the community to facilitate adoption of SDN.

Building a model for incremental replacement of the legacy network with SDN components is the right approach for production and prototype deployment of SDN. Similarly, small pieces of the network, like the Science DMZ, can be made SDN/OpenFlow compatible while providing an area for development of understanding, tools and network best practices.

There was debate around the importance of legacy Operations, Administration, and Management methods and mechanisms and if these should be supported by the new paradigm.

Moving to Multi-Domain SDNs

Most software-defined networks today are built under controlled single-user domains. Investment is needed to leverage SDN to build general-purpose IP networks that can interface with standard end-hosts and devices people use across the Internet today.

Multi-domain SDN research is needed and is not currently a priority in the marketplace. SDXs could be an approach to tackle this need.

Enabling Application Use of SDNs

Network virtualization and management of the underlay network in support of that is a critical area to explore.

The current programmability of APIs is insufficiently abstracted to allow applications to interact effectively with the network. More effort needs to be put into building the right abstractions including for monitoring and control to not only enable application programmability but also to allow network operators/engineers to manage that interaction.

Flexible description of application profiles is needed to handle the different kinds of applications since commercial application characteristics may vary significantly. Not all applications need to be network-aware or be able to program the network. These profiles need to be developed jointly by application and network engineers and can be application-specific.

Managing the underlay network: the physical network with multiple layers (like optical, Ethernet, and wireless) across different physical media needs to be exposed to the higher layers that enable applications to build flexible overlay networks over the current physical infrastructure

The models for interaction of applications and the network, including trust and security, need to be explored in order to increase the value of programmability and creation of SDN-aware applications. Defining an ecosystem for applications, that includes the interaction between storage, compute and network, needs to be addressed.

How do we get there?

The community representing the applications and networks representative of the Government and R&E community should take an active role in defining requirements of the underlying network operating system layer from their applications perspective. This will help vendors build solutions and APIs that match the requirements.

Creation of multiple SDNs connected with each other, run by operational network engineers and providing an environment where new tools can be deployed, vetted and improved will be critical in moving this new technology from the research stage into wider production. The lack of multiple, open, operational SDNs has limited the growth of practical tools, standards and best practices with an operational feedback loop.

Capable, experimental platforms are needed that could be based on open- hardware and software to enable network operations and application engineers to experiment with the various SDN concepts so they can learn and adopt the new paradigm of programmability.

This ecosystem should have participation and representation from network, and software engineers so the participants can build something and learn quickly from that experience.

5.3 Security

SDN (Inter-domain SDN) is vulnerable to global or large-scale attacks on its control plane. The Internet, through its decentralized architecture, mostly limits the success of attacks to a local or smaller scale. In-scope for SDN security research investigation are the architectural features for the self-limiting of global attacks, or providing resilience against these attacks. Inherent in this line of reasoning is the assumption that some attacks will succeed locally, whether due to insider threat or vulnerabilities in specific implementations of SDN elements such as switches or controllers.

Security of a Single SDN

Here we focus on security issues of a single-domain SDN: a network or collection of networks that all trust a common security token associated with a single administration for authentication, authorization and session management. In the simplest single-domain SDN, a uniform security policy and enforcement mechanism constrains access, external use, and adversaries attacking all network devices and controllers that are part of the SDN.

Trust models provide the security behavior or expectations of the information flows, the network and of each of the participating elements: the switches, the controller, and the

connectivity fabric. In particular:

- What is the trust relationship between switches and the controller;
- How is trust established as parts of the network are (re)booted;
- Should a switch trust all commands it receives from a controller?
- Should the controller trust information from the switches?
- What can a controller assume about impersonation of one switch as another or as a non-switch entity as a switch?
- To what degree of certainty do controllers and switches need to authenticate and authorize one another?

The research issues are in the details of what authentication (mechanisms), authorization (mechanisms), and policies to define and use (who can do what to whom). We observe that a trust model must cover at least all of the elements, and the decisions made about each will determine a set of additional threats that need corresponding mitigation.

The programmability of an SDN controller blurs the distinction between network and application security. This may permit more holistic protection and potentially earlier detection of threats. In that setting, the controller needs information that is currently typically unavailable at the network level, about application-level principles. Research should help define explicit mechanisms and interfaces by which application-level protocols and services may expose information about principals, suitable for authentication, authorization and policy enforcement without compromising the overall security of the network.

Once an SDN is in service, it will require on-the-fly upgrading. This is a broad topic encompassing software, hardware, topology, protocol, and media; requiring both static and dynamic support. Because upgrades cannot be performed atomically, there are in principle two (or more) networks overlaid atop one another, yet traffic needs to be treated in a semantically sensible manner including, for example:

- from an individual packet's perspective: it should experience consistent treatment from the network.
- at a higher level, this same guarantee is desirable for not only a packet but for an entire logical flow.

Finally, SDN controllers are highly likely to run a collection of apps representing differing functionality – everything from existing network protocols and services to novel, custom applications that exploit the network visibility provided by SDNs. This confronts us with the interaction between (obviously) conflicting or even (actively) competing applications. What notions of modularity will enable us to safely compose these applications in predictable ways?

Security of Multiple, Decentralized, Interdomain SDNs

The future deployed SDN infrastructure must deal with real-world complexities:

1) the presence of a large number of distinct autonomous domains,
2) an explicit lack of trust between all, "the entirety", of the autonomous domains.

Thus:

1) SDN security architecture and mechanisms must be designed for eventual scalability;
2) SDN security policies must support pairwise and explicit negotiation of security services – information visibility, mutual (pairwise) policy enforcement, while enabling and supporting security services for end-users that span the entire SDN infrastructure ("SDN-wide security objectives").

This is in direct opposition to a simpler model, in which all autonomous domains, each running a single local SDN in an inter-domain SDN eco-system, sign-up and agree to a common, universal set of security policies and mechanisms.

The following list identifies key areas for development and research investigation in support of creation and operation of a production-level large-scale Inter-domain SDN.

- *Abstraction of Shared Information*: Each autonomous SDN domain should maintain an adequate view of configuration information and the operational state of all the elements (switches, controllers, administrative systems and databases) under its control if it is to ensure security requirements. However, for both technical scalability and practical security posture, autonomous SDNs must limit the sharing of information concerning internal elements with (i) immediate peer SDNs and (ii) the global collection of indirectly connected SDNs. Research investigations are needed to understand this complex balance.

- *Inter-domain SDN (multiple, decentralized) properties*: Research is needed to distinguish the properties that can be delivered to end-users of an inter-domain SDN from multiple constituent SDNs operating in a decentralized manner. Such research spans topics including appropriate definitions and enforcement of service level agreements. Specific properties such as QoS (interpreted very broadly as any network forwarding characteristics beyond "best effort"), and resiliency of the provisioned SDN services (beyond packet forwarding, to include programmable actions) need to be worked out and deployed in a practical setting for inter-domain SDNs. Also, the degree to which control services are exposed and secured must be investigated. This specifically includes the ability to insert or remove various application-specific "flowspace" handlers at various SDNs (originating SDN, transit SDN, destination SDN, etc.), as well as the resource limits or other security properties imposed by policy on these control service interfaces.

- *Authorization*: Each local SDN is presumed to have a trust model for local authentication, authorization, and policy enforcement. These trust models may vary from very limited to extremely sophisticated depending on the size and

requirements of each autonomous SDN. For the operational Inter-domain SDN, authorization solutions require:

- A pragmatic definition and adoption of common vocabulary for use in defining and sharing authorization policy;
- Establishment of global terms that are authoritative when used, while permitting subsets of the inter-domain SDN to introduce additional terms with specific shared meaning; and
- A common understanding of what defines a principal.

On this last point, each type of principal referenced by inter-domain SDN policy must support the export of the authentication information proving that a specific principle is initiating an action, so that actions may be securely bound to principals across the requisite portion of the inter-domain SDN.

- *Consistency Models*: Research is required to investigate how explicitly to exchange security information, or to expose security interfaces, between and among applications and elements of the inter-domain SDN. This also may require research into how to securely maintain the consistency between applications and controllers, as well as applications and overall network state. (For example, what application flows are understood and securely handled by elements of the SDNs across the entire inter-domain SDN?)

- *Peering autonomous system*: Production inter-domain SDNs may use current autonomous system peering architectures (such as IXPs), to peer SDNs, as has already been demonstrated in limited use. For such peered SDNs, issues arise as to how to securely peer the SDN-aware applications that span such an inter-domain SDN. This topic of investigation may be replication of some of the previous topics, as applications already may incorporate their own security notions, and the challenge here may be to integrate smoothly such notions as transparently as possible across multiple SDNs, which may be late to the game in establishing their own secure inter-domain SDN relationship. (The applications may precede the SDNs in terms of security interoperation.)

Design Principles for Secure Control Plane

The security of an SDN network is dependent on the control plane. This dependency suggests that there should be a set of design principles and patterns for building a secure control plane: its northbound APIs and exposure of information, frameworks for securing the ecosystem of controller applications, monitoring and anomaly detection, the complexity of multiple tenancy and levels of acceptable interference between flows, the guarantees of achievable non-interference if resource perturbation is avoided, how to implement redundancy, and how coordination can be accomplished within the flow space (for example between monitoring flows.)

Classification of properties and separation

SDN networks introduce new networking abstractions and properties and these offer ways to specify requirements and create policies. For example, one can define isolation properties within an SDN such that information is not flowing between two subnetworks of the network. Flows within the network become an important first class abstraction and lead to research classifying the properties of SDN networks that offers new ways to control and secure these systems.

The separation of the control and data plane permits the use of mitigated interference in allocating bandwidth and resources. In a similar manner, the notion of admission, quality of service, off-line analysis, and revocation control of data and control flow packets within an SDN permits a classification of validated flows. In general, classifying properties of flows as first class abstractions to capture their static and dynamic behavior and consequent security properties and concerns is of research interest because of the programmatic nature of SDN.

Uses of SDN for Security

Allowing SDN-enabled security provisions to exploit the new abstractions and programmability of the networks introduces the possibilities of research to use these technologies in security policies and security applications. Such research might examine secure updates for applications, flexible intrusion detection, and flexible reaction and provisioning.

6 Appendix: Workshop Participants and Contributors

First Name	Last Name	Organization
Frank	Acker	NSA
George	Adams	Purdue University
Aditya	Akella	University of Wisconsin-Madison
Ali	Al-Shabibi	Open Networking Lab
Josh	Bailey	Google
Edward	Balas	Indiana University
Ilya	Baldin	RENCI/UNC Chapel Hill
Nick	Bastin	Barnstormer Softworks
Gregory	Bell	ESnet
Joseph	Berthold	Ciena
Bob	Bonneau	AFOSR
Eric	Boyd	Internet2
Joe	Breen	University of Utah
Doug	Butler	
Roy	Campbell	University of Illinois
Richard	Carlson	DOE-SC/ASCR
John	Carter	IBM Research - Austin
Russell	Clark	Georgia Institute of Technology
Steve	Corbato	University of Utah
Eli	Dart	ESnet
Vince	Dattoria	DOE-SC/ASCR
Narayan	Desai	Argonne National Lab
Chip	Elliott	BBN
Joseph	Evans	University of Kansas
Nick	Feamster	Georgia Institute of Technology
Darleen	Fisher	National Science Foundation
Nate	Foster	Cornell University
Luke	Fowler	Indiana University
Cliff	Frost	CENIC

Navid	Ghazisaidi	Verizon
Chuck	Girt	OneCommunity
Matthew	Goodman	DARPA
Paola	Grosso	University of Amsterdam
Chin	Guok	ESnet
Deniz	Gurkan	University of Houston
Hwa-Jung	Han	Verizon
Ron	Hutchins	Georgia Institute of Technology
Julio	Ibarra	Florida International University
Eiji	Kawai	NICT
Stephen	Kent	BBN
Shriram	Krishnamurthi	Brown University
Keith	Landgraf	DoD
Larry	Landweber	University of Wisconsin
Michael	Langdon	Juniper Networks
Bryan	Larish	Georgia Institute of Technology
Tom	Lehman	Mid-Atlantic Crossroads (MAX)
Bryan	Lyles	National Science Foundation
Iara	Machado	RNP (Rede Nacional de Pesquisa)
Joe	Mambretti	International Center for Advanced Internet Research Northwestern University
Scott	McNown	NSA/R2
Linden	Mercer	Naval Research Lab (ARL/PSU)
David	Meyer	SP CTO
Grant	Miller	NITRD
Vinod	Mishra	Army Research Lab
Inder	Monga	ESnet
Anita	Nikolich	National Science Foundation
Ping	Pan	Infinera
Eric	Pouyoul	ESnet
Nagi	Rao	Oak Ridge National Laboratory
Jennifer	Rexford	Princeton University
Glenn	Ricart	US Ignite

Christian Esteve	Rothenberg	University of Campinas
Daniel	Schmiedt	Clemson University
Stephen	Schwab	USC Information Sciences Institute
Syed	Shah	OSD
William	Snow	Open Networking Lab
Lorenzo	Vicisano	Google
Robert	Vietzke	Internet2
Ann	Von Lehmen	Applied Communication Services
Steven	Wallace	Indiana University
Robert	Walter	DARPA Strategic Technology Office
James	Wanderer	Google
Kuang-Ching	Wang	Clemson University
Gerhard	Wieser	Cisco Systems
John	Wilbanks	NCO
Linda	Winkler	Argonne National Laboratory
Hagen	Woesner	BISDN GmbH
Michael	Zink	University of Massachusetts Amherst

7 Acknowledgements

ESnet is funded by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research (ASCR). Vince Dattoria is the ESnet Program Manager.

ESnet is operated by Lawrence Berkeley National Laboratory, which is operated by the University of California for the U.S. Department of Energy under contract DE-AC02-05CH11231.

This work was supported by the Directors of the DOE Office of Science, Office of Advanced Scientific Computing Research, Facilities Division.

This workshop was also supported in part by the National Science Foundation under Grant No. 0936815.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding agencies.