Advanced Layer 2 Service Network Risk Mitigation

I. Introduction

The Advanced Layer 2 Service Network is intended to be a production-quality, 100-Gigabit, Software-Defined Network (SDN). This document will describe the risks associated with the migration of the current infrastructure to the Advanced Layer 2 Service Network and discuss strategies to mitigate those risks.

Until recently, most network operators have thought of SDNs, implemented using the OpenFlow protocol, as experimental networks to serve as breakable test-beds for network researchers. Over the past few months there have been announcements of production networks built using the SDN architecture, most notably by Google. However, there is limited information available publicly about these deployments and therefore no roadmap for building such a network and limited products or open-source software available.

While this creates significant risk to large-scale implementers, the Internet2 community has a unique opportunity to be a leader and innovator in the area of building high-bandwidth, programmable networks.

Therefore, it is the opinion of the NTAC that:

- Quality of the existing Layer-2 and Layer-3 network services the community relies on be preserved to match the quality experienced on the current infrastructure.
- It should be well-stated that the goal of the Advanced Layer 2 Services Network is to be a *production quality* software-defined network.
- Best practices for building production SDNs do not yet exist, so it should be recognized that for a period of time the Advanced Layer 2 Services Network will be operating in a pre-production state as such practices are identified and implemented.

With an articulated goal of production services, a brief survey of the current environment includes the following challenges:

- The SDN standards being developed by the Open Network Foundation (ONF) are still very new and have been evolving rapidly.
- The ecosystem of SDN products and open-source software is also nascent and evolving rapidly.
- Network engineers at campuses and regional network have very limited experience with SDN technology.
- There is no standard API for SDN controllers which makes it challenging to develop software that can be easily deployed across many different, autonomous networks.
- There are mechanisms in place for compliance testing. However, there is a lack of production testing, scalability and stability remains unproven on production systems.
- Funding sources are limited and member's technology refresh timeline may not align closely with this opportunity.

All of these challenges can be addressed and overcome, but it will require careful planning, full transparency, and robust testing. One cannot understate that the success of this endeavor, to a great extent, rides on the "buy-in" and confidence of a number of community constituencies. Articulating a prime directive of "do no harm" to existing services will go a long way toward easing some of the trepidation being articulated at this time.

The technical risks associated with the Advanced Layer 2 Services Network can be divided into two categories: risks associated with the SDN technology itself and risks associated specifically with the Advanced Layer 2 Services Network. In the following sections we will enumerate those risks and suggest steps for mitigating those risks.

II. SDN Technology Risks

a. Immaturity of Technology and Products

Similar to other early-stage technologies the Internet2 community has deployed in the past, both the technology and the products needed to deploy an SDN/OpenFlow network are immature. There are many parallels to the deployment of inter-domain multicast within the Internet2 community in the late 1990's and early 2000's. The Internet2 community should draw on those experiences as we plan for the Advanced Layer 2 Services Network deployment. Some of the programs that have worked in the past include:

Hands-On Workshops: The Internet2 hands-on workshops for IPv6 and Multicast were both popular and successful. Many campus and network engineers received valuable instruction on how to deploy these technologies. They were able to return to their campuses and start deploying these technologies.

Experience Sharing: When they returned to their campuses, they often had questions and were able to ask for help either through the working group email lists or direct email to instructors or fellow attendees. In this age of Web 2.0 applications, there are likely other avenues we could leverage to improve this experience.

Common Deployment Models: It's important to have a clear set of steps or phases for deploying a new technology. With IPv6 and Multicast, there were clear instructions on what technologies to use, how to configure them and how to phase in a deployment. Similar guidance should be provided with SDN/OpenFlow.

Coordinated Feature Requests to Vendors: It's important that the community speak as one, unified voice with the networking vendors. There are many new features the vendors could be implementing that could be more or less valuable to what this community is trying to receive.

b. Splintering of Technology into Proprietary Solutions

The development of standards typically proceeds at a slower pace than market demand for solutions. This leaves a gap in which vendors can differentiate their product by delivering functionality before it's standardized. This can be helpful, but it's important that these advancements are rolled into the standard. In the case of SDN/OpenFlow, this risk is increased by the fact that many people consider this a disruptive technology. Suggestions for mitigating this risk include:

- Requiring vendors to provide a path to standardization for any OpenFlow extensions that the Innovation Platform network utilizes, and
- Requiring vendors to engage in compliance and interoperability testing.

As another example of "speaking with one voice," the community should develop and maintain a simple implementation agreement for OpenFlow products, essentially a checklist of mandatory and optional functions in the standard that the community agrees are required. Publishing such a document, making vendors aware of it, and referencing it during procurement could help nudge vendors towards supporting commonly required features. It should be noted that this type of agreement is a complement to interoperability testing, not a substitute for it.

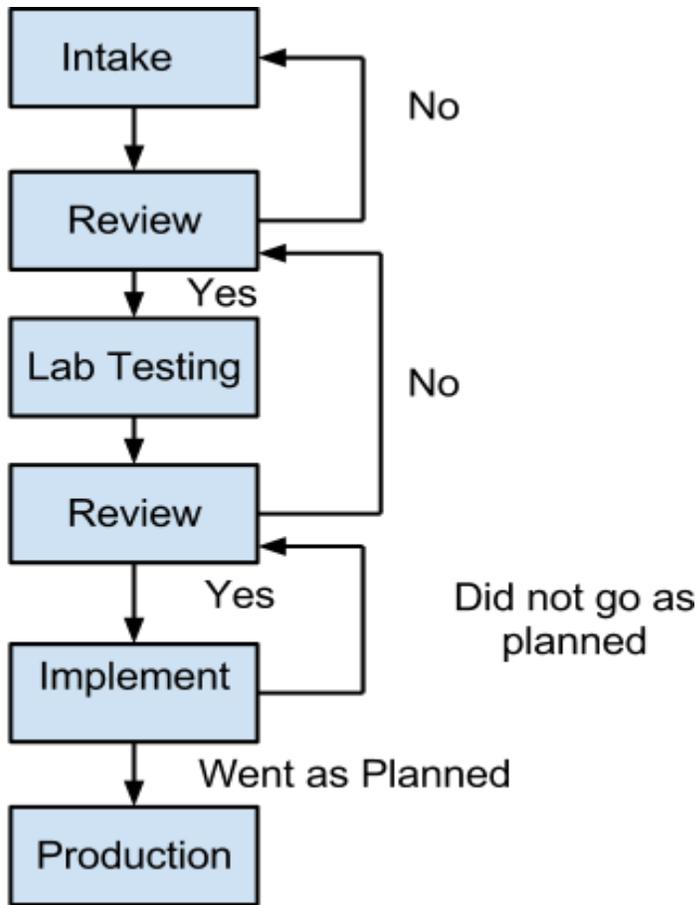## III. Advanced Layer 2 Services Network Risks

a. Balancing Support for Production Services with Support for Experimentation

Perhaps the key risk for the Advanced Layer 2 Services Network will be the balance between supporting production services and experimentation. Internet2 has stated that this network will support the GENI API to allow researchers to run experiments inside "slices" of the network. Additionally, it is the NTAC's understanding that Internet2 plans to run a production layer-2 service on the same set of switches. In order to be successful, this will require very careful planning and preparation. A requirement for success will be an ability to adapt nimbly to an evolving set of requirements and circumstances. Rather than being specifically prescriptive, it is our intention to articulate a framework incorporating aspects of feature request/evaluation, change management and peer review.

The process for review should include at least the following components:

1. All requests and review should be in the public domain; the entire process should be transparent and easily accessible to the entire community.
2. There should be a well articulated process for interested parties to engage the mechanisms for making a request.
3. All features and requests should be vetted for technical viability (presumably via the InCentre lab at IU). We do not want to cause any damage to production resources.
4. There should be a small, independent advisory panel that would work with Internet2 staff to review the request for stability and interoperability prior to engaging resources for viability testing, as well as after testing to make a judgment on the lab's recommendation.

Proposed Workflow

Intake → Review → Lab Testing → Review → Implement → Production

```
┌──────────────┐
│    Intake    │◄──────┐
└──────┬───────┘       │  No
       │               │
       ▼               │
┌──────────────┐       │
│    Review    │◄──────┤
└──────┬───────┘       │
       │ Yes           │  No
       ▼               │
┌──────────────┐       │
│  Lab Testing │       │
└──────┬───────┘       │
       │               │
       ▼               │
┌──────────────┐       │
│    Review    │◄──────┤
└──────┬───────┘       │  Did not go as
       │ Yes           │  planned
       ▼               │
┌──────────────┐       │
│  Implement   │───────┘
└──────┬───────┘
       │ Went as Planned
       ▼
┌──────────────┐
│  Production  │
└──────────────┘
```

IV. Conclusion

There is no doubt that we are witnessing a major paradigm shift in the world of data networking, the likes of which has perhaps not been seen since the original evolution of the Internet. With the implementation of the Advanced Layer 2 Services Network, the Internet2 has a unique opportunity to help usher the world into an exciting realm of production-quality, programmable, adaptive, intelligent, and high-performance networking. As with any innovative venture, there is significant risk that is best mitigated by transparent, open processes; community participation and peer review; open development of best practices; and an honest admission that we are breaking new ground and are working together as a community to determine exactly what those best practices are.