Network Development and Deployment Initiative (NDDI) and Open Science, Scholarship, and Services Exchange (OS^{3}E) Frequently Asked Questions (FAQ)

OVERVIEW: NDDI

1) What is the Network Development and Deployment Initiative (NDDI) ?

The Network Development and Deployment Initiative is a collaboration between Internet2, Indiana University and the Clean Slate Program at Stanford University to develop and prototype a nationwide Software-Defined Networking platform that will support both production services and virtualized "slices" of the network substrate upon which new protocols or services can be developed. Indiana University, Internet2, and the Clean Slate Program are contributing substantial resources to the collaboration including investments in development, hardware, waves, international relationships and operating expertise.

NDDI capabilities will be developed and interconnected with links to Europe, Canada, South America and Asia, through coordinating international partners like RNP in Brazil, CANARIE in Canada, GÉANT in Europe, and JGNX in Japan, with additional service partners to be identified.

2) How does this help scientists? How is this better than what we have now?

The NDDI substrate will allow scientists who need a customized, large-scale testbed for both new applications and new protocols to work at large scale with relative ease. Individual scientists will not need to develop and operate their own networks, maintain a roadmap for changes, capacity, and service capabilities, etc. This new substrate will provide a large scale platform that will be professionally managed to support their science so that they can concentrate on the science itself.

3) What are some typical ways people might use this?

The substrate will initially allow researchers working on projects like GENI to extend their testbeds with greater flexibility and scalability than current testbeds offer. Over time, the substrate will be flexible enough to allow custom overlays designed by individual projects. By example, a collaborative of undergraduate computer science classes might use the substrate to let undergraduate classes write their own protocol stacks and test them on a global testbed.

4) What is Software-Defined Networking? What is OpenFlow?

Software-Defined Networking is an approach to networking that allows network operators to optimize the behavior of their networks by telling the network hardware what to do with packets or frames according to user defined rules. In the same way that Linux allows great flexibility and customization as an operating system, OpenFlow offers the same sort of access to control the network forwarding approaches.

5) How can my campus network become a part of NDDI? What does that mean? What is the cost? What equipment is needed?

NDDI is expected to become a national testbed on the Internet2 footprint with access points at approximately 35 locations across the US. At the simplest level, a campus can

gain access to the NDDI by asking their Internet2 connector to provision some VLANs from the campus to the nearest NDDI switch. More advanced development efforts may wish to attach dedicated waves or circuits from campus testbeds to the NDDI substrate either directly or through a regional network.

A campus can attach workstations or other devices to NDDI by extending VLANs through its Internet2 connector to the lab on the campus. If a campus wants to experiment with NDDI capabilities to develop or use their own Software-Defined Networking tools, the first step will be to select Ethernet switches that support the OpenFlow standard. A server that can act as a controller will also be required. The cost for the switches and servers should be similar to other switch purchases.

6) How can my regional network become a part of NDDI? What does that mean? What is the cost? What equipment is needed?

At the most basic level, a regional network can attach its network devices to NDDI and extend services to the campuses it serves by extending VLANs from its Internet2 connection point to the border of interested campuses. If a regional wants to experiment with NDDI capabilities to develop or use their own Software-Defined Networking tools, the first step will be to select ethernet switches that support the OpenFlow standards. A server that can act as a controller will also be required. The cost for the switches and servers should be similar to other switch purchases. Internet2 is investigating supporting the NDDI architecture as part of the current DYNES deployment.

The cost structure for the NDDI is expected to be worked out as the final design and program costs are solidified. The costs are expected to be less than existing Internet2 port charges and may be offered at no charge to Internet2 connectors who already have multiple large connections to the backbone.

7) How can my National Research and Education Network (NREN) become a part of NDDI? What does that mean? What is the cost? What equipment is needed?

Internet2 and Indiana University are both coordinating with global NREN partners to connect them with the NDDI initiative. NRENs interested in coordinating with NDDI should contact nddi@internet2.edu, which includes representatives from Indiana University and Internet2.

8) How do virtual organizations interact with the NDDI?

Internet2 and Indiana University are both coordinating with global virtual organization partners to connect them with the NDDI initiative. VO's interested in coordinating with NDDI should contact nddi@internet2.edu which includes representatives from Indiana University and Internet2.

9) I'm interested in the NDDI and related services, how do I get involved?

Individuals or organizations interested in coordinating with NDDI should contact nddi@internet2.edu which includes representatives from Indiana University and Internet2.

SUBSTRATE / SERVICES DELINEATION

1) How does the NDDI substrate differ from related services such as the Open Science, Scholarship and Services Exchange (OS³E, the VLAN service)?

The NDDI substrate is a set of network software and hardware that is developed and deployed on a national footprint by the NDDI partnership. The NDDI substrate will be capable of supporting multiple services simultaneously. The Open Science, Scholarship and Services Exchange (OS³E) is the first service that will be deployed on the NDDI substrate. An at-scale network research testbed service that allows network researchers to create "slices" for experimentation is also being planned and more services could be added in the future.

2) How does the NDDI substrate differ from potential future support of at-scale network research testbeds?

NDDI as a substrate is intended to support multiple at-scale network research projects at the same time. Internet2 and Indiana University expect that, as the tools developed under NDDI evolve, NDDI will become an important long term and evolving resource for network research.

3) How does the NDDI substrate differ from potential future support of a distributed open science exchange? How does this relate to LHCONE?

The Open Science, Scholarship and Services exchange will be a layer-2 service operated on top of the NDDI that will operate as a distributed national exchange. Any research or science project that requires a robust layer-2 infrastructure may choose to use the OS³E for their projects. NDDI will work with the LHC community to ensure that the NDDI substrate is meeting the needs of the LHC community, as articulated in the LHCONE architecture document.

4) What are the requirements for participation in each of the aforementioned services and initiatives?

Participation in the NDDI Partnership is distinct from participation in the services deployed on top of the NDDI substrate, such as the Open Science, Scholarship and Services Exchange. Each service deployed on top of the NDDI substrate could potentially have it's own policies regarding participation. The NDDI Partnership currently has two categories of partners: An Investing Partner is a partner that contributes significant resources to the NDDI project in the form of hardware, network capacity and/or software development. Investing Partners contribute to all aspects of the NDDI project, not just a specific component or location. The initial investing partners are Internet2, Indiana University and the Clean Slate Program at Stanford University. A Cooperating Partner is defined as a partner that either implements a similar infrastructure and links their infrastructure to the NDDI substrate or contributes to specific components or aspects of the NDDI project.

POLICY

1) Who can connect to the NDDI? What is the AUP for NDDI?

Internet2 and Indiana University will work with the appropriate governance bodies to develop an operational and sustaining set of policies for the NDDI. It is the intent of the NDDI team to make the substrate widely available for network research and innovative new services that require access to the substrate. Operational and sustainability policies

will need to be developed to support overall NDDI stability and any ancillary controller support at the NDDI's start. As the substrate evolves and becomes more robust, it is hoped the policies and access will become increasingly flexible and open.

2) Who can connect to the OS³E? What are the acceptable use polices for OS³E?

The Open Science, Scholarship and Services Exchange will be operated as a distributed national exchange. Any organization that pays the port fees may connect to the OS³E and may use it for any legal purpose. As a layer-2 open exchange network, a organization that can connect to a port on the exchange may connect to any other party on the exchange who agrees to accept a connection from them.

3) What advisory and governance processes will be used for the NDDI substrate?

Internet2 will continue to engage with the R&E community through its normal governance processes with respect to the NDDI substrate, OS³E, and other ensuing services.

COSTS

1) What are costs related to participation in NDDI? In OS³E? Who gets charged?

A sustaining cost structure has not yet been developed, however the expectation is that the NDDI substrate itself will be available to Internet2 members with network participation agreements. Fees for the OS³E have also not yet been set, however it is expected that the OS³E which runs on top of the NDDI will be open to any party who pays the port fees. As an open exchange, there will not be requirements on the OS³E directly connected to Internet2 membership.

2) What is the charge for a connection? How much will a port cost?

Fees will be set during the Summer of 2011 as the final costs to operate the NDDI and OS³E are better understood. Individuals interested in participating in discussions about the fee structure should contact nddi@internet2.edu.

3) What is the fee associated with usage?

Fees will be set during the Summer of 2011 as the final costs to operate the NDDI and OS³E are better understood. Initial thoughts are there will not be usage charges. Individuals interested in participating in discussions about the fee structure should contact nddi@internet2.edu.

4) Is there an extra cost for bandwidth protection between connections?

Fees will be set during the Summer of 2011 as the final costs to operate the NDDI and OS³E are better understood. Initial thoughts are there will not be usage charges. Individuals interested in participating in discussions about the fee structure should contact nddi@internet2.edu.

TECHNICAL DETAILS

1) What type of networking is provided by the NDDI substrate? To the related services?

The NDDI substrate provides a general-purpose Layer-2 platform that can support

multiple instances of networking services atop a common network footprint. While the end-user services may appear similar in performance and appearances to a traditional VLAN-based set of services, the underlying network is very different from traditional Layer-2 network and is much more powerful and versatile. As the OpenFlow software evolves, user traffic can be routed by more than just VLAN tags, creating whole new classes of network services.

2) What is the physical network footprint of the NDDI substrate? Where will the nodes be located and when will they be installed?
The NDDI substrate footprint is evolving. We intend to have equipment at nearly all Internet2 optical add/drop facilities in addition to several off-network locations, for a total of 30-40 nodes. The network will ultimately scale up to this level as the technology matures and the use cases become more apparent. The first set of nodes is expected to be installed in the Fall of 2011.

3) Is the NDDI substrate and the OS^{3}E a fully supported service with 24x7 monitoring?

Yes, the NDDI substrate and the OS^{3}E will be fully supported services at the Internet2 NOC. Both will have 24x7x365 monitoring and engineering support. The substrate and Layer 2 service will be integrated into the suite of available tools, so users will be able to view usage reports, outage statistics, and real-time analytics.

4) How will the NDDI substrate address spanning tree issues?

The NDDI substrate will take advantage of the capabilities of OpenFlow to address issues involved in the operation of large layer2 infrastructures. OpenFlow not only allows for the creation of Spanning Tree or similar mechanism for loop prevention but also the use of multi-path techniques to maximise flexibility and throughput.

5) How much bandwidth will be provisioned between each node? What is the augmentation strategy?

The initial inter-node bandwidth will be 2 10 Gbps circuits, though the NDDI Partners intends this to grow over time as the OpenFlow-compliant hardware base evolves toward 40G and 100G port availability. The NDDI Partners will conduct periodic reviews of bandwidth consumption to maintain adequate headroom for users.

6) What are the connection sizes? 1GigE? 10GigE? 40GigE? 100GigE?

Users of NDDI substrate services will initially be able to connect at 1GigE and 10GigE speeds, though the NDDI Partners intend to widen these options as the OpenFlow-compliant hardware base evolves toward 40G and 100G port availability.

7) Is there bandwidth protection by default or available?

Bandwidth protection is on the roadmap and an ultimate requirement. Availability will be determined by the selected NDDI substrate hardware vendor and software development progress.

8) What level of stability should participants expect?

The NDDI substrate represents one of the most forward-looking advanced infrastructure investments in the past decade. As such, services atop the NDDI substrate should anticipate a low level of churn as software updates are pushed out and improved. However, the NDDI Substrate will be treated as production-level effort that is fully monitored and reported upon.

TECHNICAL DETAILS: OS^{3}E (VLAN SERVICE)
1) What latency protections will be possible?

Users will be able to select a path which minimizes latency, if desired. The VLAN Service will allow users to select the shortest available path automatically. No hard latency protections are expected to be included.

2) Will VLAN IDs be necessary?

Yes, the Layer 2 service will use VLAN IDs as the primary identifier. We expect that you will be able to use any VLAN ID on any edge port, and that VLAN IDs will be transparently translated across the backbone, as necessary, to avoid conflicts.

3) Will VLAN tags be preserved?

Yes**.**

4) Will Ethernet CoS information be preserved?

Yes.

5) If this NDDI substrate is being built using Ethernet switches, how will it avoid the spanning-tree convergence problems sometimes found in modern ethernet networks?

TECHNICAL DETAILS: Software-Defined Networking / OpenFlow

1) What is Software-Defined Networking?

Software-Defined Networking is an emerging architecture that provides an open interface to control forwarding decisions of network devices. Using technologies such as OpenFlow, the behavior of network devices can be dynamically controlled, allowing innovative new protocols and approaches to be explored, tested, and deployed quickly.

2) What is OpenFlow? Why OpenFlow?

OpenFlow is a protocol which enables a Software-Defined Networking architecture. OpenFlow allows any piece of software which implements the OpenFlow protocol (a "controller") to direct forwarding decisions of OpenFlow-compliant network devices. This is a accomplished through a set of “flow rules”, which the controller uses to inform each network device how to forward frames/packets. OpenFlow provides a standard interface for network engineers or application developers to control data flows by creating rules like "for any packet which comes in interface 1 with destination IP address 1.2.3.4 and TCP port 80, send this packet out port 2." OpenFlow is the leading technology enabling Software-Defined Networking, and The Open Networking Foundation was recently formed to continue OpenFlow standardization. We believe that building on top of an

OpenFlow substrate will allow a great deal of flexibility in supporting a variety of different services on one physical platform in a very cost effective manner.

3) Is OpenFlow mature enough?

OpenFlow has been rapidly gaining momentum. OpenFlow support has been announced by many major network equipment vendors, and several commercial and non-commercial OpenFlow control platforms have been developed. Many R&E community members have OpenFlow testbed deployments, many of which are expanding. While the standard is still evolving, we believe that OpenFlow has reached a level of maturity that can support a production service, and that we will be well positioned to enhance services as the OpenFlow standard and implementations continue to evolve.