Connecting to the GENI Network
18 August 2010

Brandon Williams
Indiana University

Contents

Executive Summary

The Global Environment for Network Innovations (GENI), a project sponsored by the National Science Foundation, is a virtual laboratory at the frontiers of network science and engineering for exploring future internets at scale. There is currently no public document that explains how international networks can connect to GENI. This document explains US R/E networks providing GENI connectivity and how to connect to them and ultimately to GENI.

Important things to know when connecting to GENI:

a. Internet2 and National LambdaRail (NLR) are the high speed research and education networks that currently provide the backbone network for GENI. They will be described and then you will need to know how to connect to them.
b. GENI will be explained because it is the network to which you will connect.
c. Explaining how to connect to GENI is the ultimate goal of this paper.

Introduction

The purpose of this paper is to explain the steps required for an international network to connect with the Global Environment for Network Innovations (GENI). I will describe three networks: Internet2, NLR, and GENI. I will explain what each network is, characteristics of each network, and how to connect to each network. I will first describe Internet2, then I will describe National LambdaRail (NLR), and finally I will discuss GENI. I am explaining what Internet2 and NLR are because they are currently the backbone networks for GENI.

Internet 2

This section explains:

1. The Internet2 organization
2. Internet2 initiatives
3. Other Internet2 details
4. The Internet2 Network
5. How to connect to Internet2

About Internet2
The Internet2 organization was created in 1996. It is composed of research and education communities. The Internet2 Network is a high speed network that research and education communities can connect to. Internet2's purpose is not to replace the commercial Internet; instead, it is providing resources where new technologies can be created. Internet2 members create open-source tools that enhance the usefulness of high-performance networking. These tools are free. The tools give existing network infrastructures more capabilities. These new technologies will eventually become part of the broader Internet. Many technologies that were previously made in the Internet2 community are part of today's Internet.

Teaching and learning
Internet2 has multiple initiatives that it wishes to accomplish. The teaching and learning initiative helps instructors create new teaching styles, makes learning easier for students, and removes geographical boundary restrictions for universities. Universities can use Internet2 for delivering live videoconferencing lectures to students, downloading 3D images requiring high-bandwidth transport at a minimum of 35 Mbps, and control a microscope by using the commercial Internet.

The Arts and Humanities initiative allows video archive storage, global performances, and easy communication among different languages. Internet2 enables the Ethnomusicological Video for Instruction and Analysis (EVIA) Project to store 150 hours of digital video that researchers and teachers can access. An instructor and student with different languages can communicate with each other, and musicians are able have world-wide performances using Internet2.

The K20 initiative gives learners the opportunity to interact with professionals. The learner can write songs with a California musician, learn Spanish with a native speaker, observe someone that is flying a jet, or ask an astronaut what their greatest fear in space is.

The Science and Engineering initiative gives scientists the opportunity to do global projects quickly and easily. Many scientific projects require technology needs that only Internet2 can provide. Scientists need to let Internet2 know about their technology needs.

Other details
The Internet2 headquarters is in Ann Arbor, Michigan. Internet2 is composed of over 330 member institutions—leading universities, corporations, government research agencies, and not-for-profit networking organizations—, 60,000 institutions in the U.S., and international networking partners from more than 50 countries. Engineering, Monitoring, and Management are supported by the Global Research Network Operations Center (GlobalNOC) at Indiana University.

Technical characteristics of the Internet2 Network
The Internet2 NOC and Internet2 connectors are the groups that will accomplish Internet2's jumbo frame goal. Large IP frames, also called "jumbo grams" or "jumbo frames," are recommended for Internet2 to increase its performance over long distances. A jumbo frame is a large packet that carries 9,000 bytes to 16,000 bytes. A packet holds data that is sent through a network. In general the limiting factor in frame size is the IP MTU (Media Transmission Unit, sometimes also called Maximum Transmission Unit), which is 1,500 bytes by default over Ethernet.

The Internet2 recommendation is to have an end-to-end IP MTU of at least 9,000 bytes today, which will be increased in time. There are multiple things that need to be done for this to happen. 9KB MTUs need to be set throughout the Internet2 backbone, which has been completed. The Internet2 NOC and its connectors have to work together to set 9KB MTUs where they peer with each other. This has not been finished. The Internet2 institutions set 9KB MTUs within their networks, at least to the high-performance end-stations. This has not yet been done.

The Internet2 NOC helps members when they need help with a problem, but the NOC is not responsible for maintaining the optical fiber backbone. Internet2's nationwide 13,500 miles of dedicated optical fiber backbone is monitored and maintained by Level 3 Communications, a leading telecommunications carrier. Internet2 has complete control of the entire network infrastructure, but does not have to pay for the underlying fiber to be fixed when it breaks. Internet2 can implement faster services or new technologies when needed, independent of the carrier.

The Internet2 IP Network supports IPv4, IPv6, scalable multicast, and other advanced networking protocols. An IP network is a computer network made up of devices that are used for an internet protocol. IPv4 is what the current commercial internet uses. The main difference between IPv4 and IPv6 is the number of addresses. IP addresses are very important for understanding the difference between IPv4 and IPv6. IPv4 uses 32-bit addresses, where IPv6 uses 128-bit addresses.

Every computer that communicates over the Internet is assigned an individual IP address that identifies the device and distinguishes it from other computers on the Internet. The physical address is the MAC address of the adapter chosen by the program. The IP address is the logical address assigned to your connection by your ISP or network administrator. IPv4 addresses have 4 numbers that are separated by decimal points.

The IP address 168.212.226.204 is an example of an IP address. 168.212.226.204 in binary form is 10101000.11010100.11100010.11001100. The binary number is important because it determines which class of network the IP address belongs to. An IP address consists of two parts: one that identifies the network; and one that identifies the node, or host. A node is any device connected to a computer network. Nodes can be as complex as computers or as simple as cell phones. The class of the address determines what part identifies the network address and what part identifies the node address.

The first 8 bits identify the network and the remaining 24 bits indicate the host within the network in Class A Networks. An example of a Class A IP address is 102.168.212.226, where "102" identifies the network and "168.212.226" identifies the host on that network. The first 16 bits identify the network and the remaining 16 bits indicate the host within the network in Class B Networks. The first 24 bits identify the network and the remaining 8 bits indicate the host within the network in Class C Networks. Binary addresses start with 1110 in Class D Networks, therefore the decimal number can be anywhere from 224 to 239. Class D networks are used to support multicasting. Binary addresses start with 1111 in Class E Networks, therefore the decimal number can be anywhere from 240 to 255. Class E networks are used for experimentation. They are not used in a standard way.

Internet2 has multiple features that members can use. Internet2 has 10 gigabits per second (Gbps) of capacity along the entire network footprint. It has a future capacity that includes a potential 40Gbps and 100Gbps interface. Internet2 is a fiber optic hybrid network. All the layers are connected to each other in a fiber optic hybrid network. A packet in Internet2 can be sent at layer 3, go through Ethernet, then go through layer 3 again on its path to its destination. Internet2 has static wave provisioning in sub-wavelength increments up to multiple 10Gbps wavelengths to allow connectors to more efficiently allocate bandwidth. Internet2 has DWDM capabilities to build separate logical networks over the same fiber facility.

Internet2’s dynamic circuit network, called Internet2 ION, provides on-demand circuit provisioning across the Internet2 Network and other partner networks. Domestic users are able to connect by Ethernet, so most current campus networks easily connect to it. Participants need to connect with ION using a virtual local area network (VLAN). Users can reserve 50Mbps-10Gbps circuits. Users have the ability to reserve circuits on both Internet2 ION fabric and other dynamic circuit networks. Each connector receives 250Mbps of permanent circuit capacity. ION supports direct IDC calls from 3rd party signaling applications. Internet2 ION is free to use. There is no acceptable use policy and no membership requirements to participate.

Internet2 offers a Commercial Peering Service to its members. The Commercial Peering Service is separate from Internet2’s network, but is still operated by the Indiana University Network Operations Center. The service enables connectors to connect with the commercial Internet. There are no additional charges for using the service. There are some services that cannot be used with the Commercial Peering Service. As an example, Indiana University uses the Commercial Peering Service and pays internet service providers for services such as YouTube.

This map shows Internet2's combined infrastructure topology

How to connect to Internet2 – GigaPoPs and Exchange Points

Universities connect to Internet2 through GigaPoPs. A GigaPoP is an access point that has connections of rates of at least 1Gbps. GigaPoPs connect university local area networks (LANs) and wide area networks (WANs) to the Internet2 network. GigaPoPs use nodes to connect participants to Internet2. U.S.-based exchange points also allow non-U.S. research and education networks to connect with Internet2. Each of the following exchange points may be used to connect to Internet2.

AMPATH is an international access point in Miami, Florida. AMPATH provides an exchange point for the United States, Latin American and Caribbean research and education networks. AMPATH supports up to 10 gigabits per second Layer2 Ethernet connections, including Ethernet VLANs (virtual LANs) mapped using next-generation SONET/SDH protocols. Participants may use Packet Over SONET (POS) / Synchronous Data Hierarchy (SDH) connections. Connectors may use Asynchronous Transfer Mode (ATM) connections. AMPATH's standard interface configuration includes support for jumbo frames, as well as IPv4 and IPv6 unicast and multicast services.

Atlantic Wave is an international exchange point interconnecting the United States, Canada, Europe, and South America. Atlantic Wave connects the critical exchange points on the U.S. East Coast: MAN LAN, NGIX-East, SoX, and AMPATH. Atlantic Wave has IP peering points in New York, Washington D.C., Atlanta, Miami, and Sao Paulo. Atlantic wave has 1 gigabit per second and 10 gigabits per second Ethernet and layer 3 capabilities.

MAN LAN (Manhattan Landing) is an exchange point in New York City that provides both Layer 1 optical and Layer 2 Ethernet connections. The exchange point is a collaborative effort of Internet2, NYSERNet, the GlobalNOC at Indiana University and the IEEAF.

Pacific Wave is a joint project between the Corporation for Education Network Initiatives in California (CENIC) and the Pacific Northwest GigaPoP (PNWGP), and is operated in collaboration with the University of Southern California and the University of Washington. Pacific Wave provides a peering fabric available at the Bay Area (Sunnyvale and Palo Alto), Los Angeles (3 sites), and Seattle. Participants may connect their layer-3 configured device to a Pacific Wave switch with either Fast Ethernet, a 1GbE, or a 10GbE connection. A network switch is a small hardware device that connects multiple networks within one local area network.

StarLight is a 1GigE and 10GigE switch/router facility, and a true optical switching facility for wavelengths. StarLight provides IP-over-lambda connections. Lambda connections allow big IP flows that overload the regular IP routing level to move to the optical level, where they get better quality of service. At the same time, the IP routing level is off-loaded and can serve smaller flows better.

National LambdaRail (NLR)

This section explains:

1. The NLR organization
2. The NLR network
3. How to connect to NLR

About NLR

National LambdaRail (NLR) was formed in May 2003 and launched in September 2003. The NLR Network Operations Center (NOC) is located at Indiana University's GlobalNOC. NLR is owned by the research and education community and represents the common interests of the nation's higher education and academic research communities in achieving robust, high capacity, experimental communications services.

NLR offers services to support specific needs of the research and education community. Co-location services, cross-connections, and fiber IRUs are some of these services. NLR's co-location service can lease a rack for a member, as well as arrange for DC and/or AC power, at any of the existing NLR nodes along the NLR infrastructure. NLR can order cross-connections within an existing NLR node facility, either in bulk form or as single-pair interconnections. NLR can have dark fiber spliced from different providers at any of the existing NLR nodes along the NLR infrastructure. NLR can obtain an Indefeasible Right to Use (IRU) for dark fiber between diverse locations, including those not currently on the NLR infrastructure. An IRU is a long-term lease or temporary ownership for part of an international cable. One major international cable owner has an IRU ownership period that is granted for 25 years. An IRU is granted by the company or consortium of companies that built the cable. An IRU gives large-scale networks and Internet service providers the ability to assure their own customers of international service on a long-term basis.

NLR provides support for the Global Environment for Network Innovations (GENI). NLR provides GENI with up to 30Gbps of capacity on three different networks. At layer 2, NLR and Cisco are providing non-dedicated prototype abilities on both FrameNet and CWave with 10Gbps services for each network. At Layer 3, NLR provides 10Gbps of bandwidth on NLR's PacketNet. NLR allows any GENI researcher independent of membership affiliation to have access to PacketNet, FrameNet, and CWave. GENI researchers have regional access through the existing NLR regional infrastructure where possible. There may be regional costs and/or restrictions.

NLR also offers an equipment co-location service to GENI. The first year of the service is free and the researcher has the option to request an extension, which is subject to review. The co-location standard offering includes rack space, DC power, cooling, up to three 1GE connections, equipment installation during regular business hours, and cross connects within the NLR cage.

Asking for the equipment co-location service begins by sending a request to the NLR Experiments Support Services (ESS) at ess@nlr.net. The request should include a list of resources. NLR PoP location, bandwidth requirements, equipment power (DC), space requirements, and cooling requirements should be included in the resource request. The type of connectivity (FrameNet, PacketNet, CWave) should be given. A description of network connections and network interface media type such as fiber LR, SR, SM, MM, and copper should be part of the resource request. Connection to NLR's Cisco 6509 switches requires an appropriate SFP module. There is OOB/console/async and management access if needed.

For projects that need resources beyond NLR's donation to GENI, the project team may submit a request for support to the NLR Network Research Council (NNRC) via the NLR ESS at ess@nlr.net.

The request should include a one-to-two paragraph project description, a list of required resources with costs (if known) as described above, and an expected duration when the resources are needed beyond the standard one year. The NNRC will review the proposal, ask additional questions as needed, and decide on the level of additional support that can be offered to the project.

The GENI equipment co-location process, whether offered to GENI or through NNRC approval, will proceed as follows. The research team will confirm that it wants to proceed. An INSTALL ticket is created. NLR Layer2/3 NOC personnel will plan and manage the installation. The shipping address for the equipment and when the research team plans to ship the equipment is confirmed. Likely shipping destinations are the IU GlobalNOC or directly to the NLR PoP. The installation is scheduled. The preferred installation is sometime during regular business hours, given approximately four weeks for advance planning. Installation, testing, and debugging if necessary are all done.

Technical characteristics of the NLR network

WaveNet is a layer 1 service. You connect to NLR at layer 1 by having a device connect the NLR layer 1 network to an outside layer 1 network. There is fiber that runs underneath the ocean that allows layer 1 connection all around the world. NLR WaveNet provides point-to-point, high-capacity 10-gigabit Ethernet LAN-PHY, 40 GE LAN-PHY, or OC-192 lambdas between any two nodes on the NLR infrastructure. The NLR WaveNet service is a full-production, unprotected point-to-point wavelength. Users requiring protection may purchase a second wavelength and create protection switching at their site. NLR may be able to provide diverse routing between the primary and protection circuit.

WaveNet provides point-to-point 10Gbps OC-192 and 10G Ethernet (LAN-PHY). Synchronous Optical Networking (SONET) are standardized multiplexing protocols that transfer multiple digital bit streams over optical fiber using lasers or light-emitting diodes (LEDs). Optical Carrier (OC) transmission rates are a standardized set of specifications of transmission bandwidth for digital signals that can be carried on Synchronous Optical Networking (SONET) fiber optic networks. Transmission rates are defined by rate of the bitstream of the digital signal and are designated by hyphenation of OC and an integer value of the multiple of the basic unit of rate, OC-48 is an example. The base unit is 51.84 Mbit/s (Mbps). The speed of optical-carrier-classified lines labeled as OC-n is n × 51.84 Mbps. OC-192 is a network line with transmission speeds of up to 9953.28 Mbps.

FrameNet is a layer 2 service. FrameNet offers Ethernet-based transport services over the nationwide NLR optical infrastructure. FrameNet provides a dedicated, point-to-point Ethernet link with guaranteed bandwidth. Rates from subgigabit-per second to 10 gigabit-per-second are available. NLR offers multipoint Ethernet at layer 2 as well. Multipoint Ethernet is a service with no bandwidth guarantees that connects more than two end points specified by the user. Also, dedicated and non-dedicated Virtual Local Area Network (VLAN) services and a dynamic circuit configuration tool are offered. The Sherpa VLAN configuration tool provides guided, secure, interactive dynamic circuit configuration. It allows authorized users to provision, modify, enable, and disable dedicated or non-dedicated VLANs on FrameNet in real-time, without requiring help from the NLR NOC.

PacketNet is a layer 3 service. PacketNet is a stable, production-quality, non-interruptible, AUP-free routed IP network on a 10-gigabit Ethernet backbone that provides a range of IP-based services, including IPv4 Unicast, IPv4 Multicast, IPv6 Unicast and IPv6 Multicast. PacketNet provides a 10GE (Gigabit Ethernet) primary connection and a 1GE backup connection. The backup connection runs the opposite direction and connects to a different router.

The IP Virtual Private Network (VPN) service is a stable, production-quality, non-interruptible service using RFC2547 BGP VPNs for projects that require an overlay IP network with separation of traffic, routing, and policy from the Routed IP Service. A VPN is a computer network that is layered on top of an underlying computer network. Data travelling over the VPN is basically not visible to the underlying network traffic. The traffic within the VPN appears to the underlying network as just another traffic stream to be passed. A VPN connection can be envisioned as a "pipe within a pipe", with the outer pipe being the underlying network connection.

The NLR VPN service is available to any number of sites, using new or existing PacketNet 10-gigabit and one-gigabit connections. Projects interested in the IP VPN Service will be evaluated on a case-by-case basis to determine feasibility of IP VPN Services for the project and any related costs.

NLR offers a service to its members called Sherpa. Sherpa allows NLR members to set up vLANS between NLR members of varying bandwidth and duration, all done with no network operator intervention. Sherpa also allows “construction” of the vLAN by the requestor, or it will automatically generate the shortest-path vLAN.

NLR also offers its own TelePresence Exchange, which enables both point-to-point and multi-point TelePresence sessions between as many as 12 different physical locations and with as many as 48 simultaneous high-definition screens. Pre-set TelePresence "meeting rooms" and a common directory and dialing system for TelePresence endpoints on NLR help provide users with a straightforward, high-quality experience. TelePresence is supported by PacketNet. PacketNet is a certified Cisco TelePresence-ready backbone. Call signaling, multipoint conference capabilities, and network management equipment support TelePresence services for NLR members and participants.

NLRview and NLR’s remote hands service are other services that are specific to the needs of the research and education community. NLRview is an infrastructure deployed at NLR PacketNet. FrameNet provides the support for network measurement and other network research applications. Parts of the infrastructure are PerfSONAR compliant. NLRview has also been designed to help the NLR Layer2/3 Service Center identify and troubleshoot performance problems on the NLR backbone or between NLR and its member regional optical networks (RONs) and their customers. NLR’s remote hands service can arrange remote hands support for members using the current remote hands contracts NLR has in place.

This is a map of NLR's infrastructure

NLR Infrastructure Capabilities	
NETWORK LAYER	TECHNOLOGY
1	Optical networks
2	Switched networks
3	Routed networks
4	End-to-end transport protocols
5-7	Middleware and applications

The higher layers in a network tell the lower layers what they want to have sent. The lower layers send the information. As an example, information sent at layer 3 asks layer 2 what paths are good so that layer 3 knows how to send the information.

NLR has many features that it provides for its members. NLR has 10GE service, 40GE service, and a path to a 100GE service. NLR is 12,000 miles long coast to coast. There are points of presence in 31 cities and 21 states. NLR is the equivalent of 160 nationwide 40GE networks. NLR is the world's only 10G wide area network (WAN) cloud. A network cloud is the unpredictable part of any network through which data passes between two end points. Clouds exist because the path a packet takes between any two points in a packet-switched network can change for different packets. NLR uses dense wave division multiplexing (DWDM). A DWDM network transfers multiple signals simultaneously over one fiber optic cable. The signals are restored to their original state with erbium doped fiber amplifiers (EDFAs). In a WDM network, multiple signals are sent over a fiber by different frequencies of light. You set what light frequency you want to send and receive. You can only receive light frequencies that are the same light frequency you set to receive. A WDM network becomes a DWDM network when the number of colors or optical wavelength channels is above 20 in a WDM system.

How to connect to NLR

NLR uses some of the exchange points that Internet2 uses. PacificWave, StarLight and Manhattan LAN (MAN LAN) provide layer 1, layer 2, and layer 3 connections with NLR. You connect to WaveNet, FrameNet, or PacketNet through PacificWave, StarLight, or MAN LAN. Once connected to NLR, the services offered by NLR are available to you.

The GENI Network

<Much of this material was abstracted from various GENI documents and presentations. Pointer to these documents and presentations are available in the References section.>

This section explains:

1. The GENI organization
2. The GENI rationale
3. GENI control frameworks
4. The GENI structure

GENI overview

The Global Environment for Network Innovations (GENI), a project sponsored by the National Science Foundation, is a virtual laboratory at the frontiers of network science and engineering for exploring future internets at scale. GENI supports at-scale experimentation on shared, heterogeneous, highly instrumented infrastructure. GENI provides collaborative and exploratory environments for academia, industry, and the public to speed up major discoveries and innovation in global networks. GENI has two parts to its mission. The first part of the mission is to open the way for transformative research at the frontiers of network science and engineering. The second part of the mission is to inspire and accelerate the potential for groundbreaking innovations of significant socio-economic impact.

GENI rationale

There are multiple problems that we face in today's Internet. There are science issues because we cannot currently understand or predict the behavior of complex, large-scale networks. There are innovation problems since we currently face substantial barriers to innovation with novel architectures, services, and technologies. There are societal issues because we increasingly rely on the Internet but are unsure that we can trust its security, privacy, or resilience. GENI will perform two types of experiments to address these issues. The first types of experiments are controlled and repeated experiments that will greatly help improve our scientific understanding of complex, large scale networks. This type of experiment will allow improvements to be made in today's Internet that will help reduce the issues that the current Internet has. The second types of experiments are experiments with clean slate thinking. Clean slate thinking experiments will try things that are not like how networks operate right now. These experiments might become part of the future Internet if they are successful. The clean slate thinking experiments will help create a future Internet that reduces or eliminates problems that the current Internet has.

The origins of GENI prototyping can be traced to August 31, 2006, when the National Science Foundation Directorate for Computer and Information Science and Engineering released a major solicitation titled "Global Environment for Networking Innovations (GENI): Establishing the GENI Project Office (GPO)." BBN Technologies was one of many organizations that responded to the solicitation. In May 2007, NSF selected BBN to create and operate the GPO. Based in Cambridge MA, BBN has made its mark not only in computer networking—BBN built the ARPANET, widely known as the starting point of the today's Internet—but also in areas as diverse as informatics, sensors, and cybersecurity. GENI is being designed and built by the community, with overall project management and engineering support from staff at the GENI Project Office.

GENI is being created in a series of spirals. A spiral is a yearly project that has specific goals. Spiral 1 completed some of the overall goals for GENI. The main goal for Spiral 1 was to create a suite of experimental infrastructure in which all components (computer servers, switches, storage, sensor networks, etc) would be highly programmable, virtualizable, and federated. These goals were successfully achieved. A highly programmable infrastructure allows each component to be uniquely configured for a particular experiment. A virtualizable infrastructure enables a researcher to reserve an end-to-end virtual space across diverse heterogeneous components. An experiment's end-to-end virtual space is called a slice. Federation permits any organization with compliant infrastructure to attach that infrastructure (to federate) to GENI.

GENI is being built and operated as a federation of smaller, semi-autonomous suites of infrastructure. Multi-project clusters created in Spiral 1 are already demonstrating excellent results in linking many kinds of GENI prototypes with their control and data interfaces. Initial deployment of prototypes has been completed with compute resources and programmable network gear now running in campus and backbone networks across the United States. A steady amount of integrated demos have been performed, each showcasing increasing capability and strawman experiments, many of which span multiple prototypes.

Most of the Spiral 1 prototypes were integrated into five control plane clusters, each representing a potential architecture for the GENI design. The clusters are competing to see which type of control framework, or frameworks, will eventually be used to connect GENI components and user-level services into a single system. The five control framework clusters are PlanetLab, ProtoGENI, ORBIT, ORCA/BEN, and TIED.

GENI frameworks

PlanetLab is a cluster based on the well-known PlanetLab implementation from Princeton University, focusing on experimentation with distributed virtual machines over the Internet. PlanetLab is one of the Internet's most widely used network research testbeds. More than 1,000 researchers at top academic institutions and industrial research labs have used PlanetLab since the beginning of 2003 to develop new technologies for distributed storage, network mapping, peer-to-peer systems, and query processing. PlanetLab currently comprises 1056 nodes at 490 sites.

ProtoGENI, a cluster based on the widely used Emulab system from the University of Utah, emphasizes network control and measurement. The ProtoGENI team, led by Robert Ricci, states that “ProtoGENI is a smaller, less fancy, but functional version of both the GENI software and deployed hardware.” Because of this, the ProtoGENI team has provided GENI researchers with an invaluable framework for developing, debugging, and evaluating new systems. ProtoGENI supports experiments with computers and virtual machines, wireless networks, software-defined radios, and sensor networks.

ORBIT, from WINLAB at Rutgers University, emphasizes wireless networks and experiment workflow tools. Led by Professor Marco Gruteser, ORBIT provides both a major, open-access wireless networking testbed and OMF, a well-developed system for describing, controlling, and instrumenting research experiments across a range of technologies. The ORBIT team has also created the “GENI-enabled” WiMAX base stations that will be deployed across multiple campuses in GENI Spiral 2.

ORCA / BEN, from Duke University and RENCI, brings both a control framework organized around flexible "resource leasing" and an experimental optical network linking multiple research organizations in North Carolina. Ilia Baldine of RENCI is leading the combined effort as well as creation of the Breakable Experimental Network (BEN), with Professor Jeff Chase leading development of the Open Resource Control Architecture (ORCA). This control framework now controls several large suites of GENI infrastructure in the sensor and wireless networking fields, in addition to optical networks.

OpenFlow is an open standard that allows you to run experimental protocols in production networks (from the OpenFlow page: http://www.openflowswitch.org/). OpenFlow is used today by major universities to develop flexible “software defined networks” as an alternative to the current hardware centric networks. See the OpenFlow URL referenced above for much more detailed information.

GENI architecture details

The GENI architecture starts with prototypes. Arbitrary numbers of prototypes are grouped in clusters. A cluster is composed of two or more integrated prototypes. In general, each cluster consists of one prototype clearinghouse plus some number of prototype aggregates that the clearinghouse controls. A clearinghouse is a management framework that tracks which infrastructures and services are available, who is authorized to use them, and which are already booked for other research or scheduled downtime. An aggregate is a GENI testbed that is managed as a coherent whole—anything from a regional network to a backbone, metro wireless, or sensor network. Aggregates provide resources to the clearinghouses. Clusters will merge into federations, which could comprise backbone networks, compute clusters, wireless networks, or sensor networks. A federation is an interconnected set of resources that are independently owned and autonomously administered.

Once federations are set up, GENI researchers will run experiments over their own slice of the infrastructure. From a researcher's perspective, a slice is a collection of computing and communication resources capable of running an experiment. Slices cut across substrates—GENI's physical plant, system software, and measurement infrastructure. From an operator's perspective, slices are the primary abstraction for accounting and accountability.

Spiral 2 will create a bigger GENI infrastructure and conduct experiments from NSF awards. In October 2009, NSF announced two major new rounds of National Science Foundation funding for the GENI project. The first award supports an additional 33 academic/industrial prototyping teams, and—for all the prototypes—speeds up federation and experiments that will guide future GENI system design. The second award supports three collaborating sets of academic/industrial teams to integrate, operate, and host experiments on an end-to-end prototype GENI infrastructure built from GENI-enabled commercial hardware across 14 university campuses, linked by compatible buildouts through two U.S. national research backbones across an aggregate national footprint of 40 gigabits/second. The first group of prototypes will use the OpenFlow protocol to integrate the proprietary hardware. The second will use the WiMAX wireless protocol, and the third the ShadowNet testbed software developed by AT&T. Both national backbones will be built out with OpenFlow. In addition, Internet2 will be built out with GENI-enabled commercial routers. The new "meso-scale" (medium-scale) infrastructure will support entirely new forms of network science and engineering experimentation at a much larger scale than has previously been available.

Connecting to the GENI Network

There is no formal paper explaining exact steps that international networks need to take to connect to GENI at the present time. This document attempts to present a minimal set of steps and reference materials to aid international researchers in connecting to GENI.

1. In order to connect to the GENI network you must receive permission to connect with GENI from the GENI Project Office (GPO). This requires the following:

1.1 Contact the GPO at geni-input@geni.net and be prepared to provide documentations to explain/discuss the following issues in a GENI Connection and Work Plan (GCWP). The GCWP should be the basis for outlining how your experiment will both connect with GENI and interact with your chosen framework. At a minimum it must address the following issues:

1.2 Explain what the goal(s) for your experiment is (are).

1.3 List what partners will be involved with your experiment.

1.4 Describe why your experiment needs to establish connection to the GENI network. What are the benefits of your connection to the GENI network?

1.5 Specify the control framework you are interested in connecting to, explain why you have selected this framework, and demonstrate you have their permission/agreement to proceed ahead.

2. Ask your international service provider how they connect to the U.S. infrastructure. Do they connect to Internet2 or NLR or both? Where exactly do they connect and what is the exact connection mechanism? This must be a part of the GCWP.

3. Depending on the answers to #2 above, make a decision to connect to Internet2 or National LambdaRail (NLR) at a specific location. This must be a part of the GCWP.

3.1 Working with your international service provider, connect to Internet2 or NLR at layer 2 at one of their exchange points. At this point, your project should have IP level connectivity to the entire US R/E community and your designated GENI control framework. If this is the required level of connectivity for your project, you are FINISHED (with the network connectivity part…). If you need dedicated connectivity between your project and your target control framework, continue on to Step 4.

4. Work with Internet2 or NLR to plan for the set-up of a dynamic circuit between the connection exchange point and the control framework node you choose to work with. This process and the Internet2/NLR agreement must be made clear in the GCWP.

4.1 Implement a dynamic circuit between the node and Internet2 or NLR landing point by using either:

4.2 Use the Internet2 ION network to make a dynamic circuit on the Internet2 infrastructure.

4.3 Use the Sherpa VLAN tool to set up a dynamic circuit on the NLR infrastructure.

5. At this point, your international experiment should have established end-to-end connectivity to the desired framework in the US and you are prepared to begin your GENI experiment which should be detailed in your GCWP.

References

http://www.answers.com/topic/gigapop

http://compnetworking.about.com/od/itinformationtechnology/l/bldef_node.html

http://www.geni.net/ (definitive source of information about the GENI project)

http://globalnoc.iu.edu (source of information about current operational status of Internet2 and NLR)

http://gordoncook.net/NLR.pdf

http://icfa-scic.web.cern.ch/ICFA-SCIC

http://www.ieee.ca/canrev/canrev37/song_eng.pdf

http://www.internet2.edu/ (additional information about Internet2)

http://www.ipv6.com/articles/general/ipv6-the-next-generation-internet.htm

http://www.nlr.net (additional information about NLR)

http://noc.net.internet2.edu/i2network/maps--documentation/policy-statements.html#JumboFrames

http://www.openflowswitch.org/

http://www.pnw-gigapop.net/services/peer.html

http://sbc.webopedia.com/TERM/G/gigaPOP.htm

http://searchnetworking.techtarget.com/sDefinition/0,,sid7_gci211800,00.html

http://searchunifiedcommunications.techtarget.com/sDefinition/0,,sid186_gci213822,00.html

http://staff.psc.edu/benninge/networking/NLR/NLR_GENI_colo_info.html

http://www.technologyreview.com