

NAOPpag Monthly Meeting Notes

9/11/19, 2:00-3:00 ET

Notes

Attendees:

James Deaton, Rod Wilson, Wendy Huntoon, Brad Fleming, Michele Norin, Jim Stewart, Marc Wallman, Scott Valcourt, Celeste Anderson

Unable to Attend: Harvey Newman, Dee Childs

Staff: Rob Vietzke, Kathleen Kay, Linda Roos, Paul Howell, George Loftus, Chris Wilkinson

1. Welcome [James Deaton]

2. Discussion Items

Optical Update [Rob Vietzke]

Internet2 recently signed paperwork for replacement fiber. We will have access to a roll pair and upgraded fiber to install on most segments. Equipment is arriving and some is ready to install; though the installation timeline has been extended, we are pressing for completion of Chicago to Denver prior to Supercomputing. There is a phasing map of 11 start to finish installations and work is ongoing to develop a schedule. Completion of the schedule is impacted by Internet2's desire to avoid long outages.

Community Architects' Workshop [James Deaton/Jim Stewart]

James reported that a report on the workshop is forthcoming. Jim Stewart emphasized the amount of good information that came out of the meeting, information that is now driving decision making at local campuses.

Linda informed the group of upcoming webinars, blogs and a newsletter to keep people updated on NGI progress and developments.

Service Model [Linda Roos/George Loftus]

- George provided background on the Service Model subcommittee chaired by Dee Childs. The subcommittee has presented the proposed Platform Fee model to the principals as a group and called regionals directly to share. Rob presented the Model to the Internet2 Board of Trustees at their July meeting, at which time they discussed the proposed model and offered feedback. The Board encouraged a focus on the new value available for members in support for research, cloud academics and enterprise with NGI.

George then shared the document prepared in conjunction with the NAOP service model subcommittee re: NAOP Platform Fee proposal. The document can be found attached at the conclusion of these notes.

Some key ideas from the proposal:

- Offer bundled services with fees that will sustain network
- Take advantage of new technology- not speed for fees

- Displayed NGI Service Model, including what is included in the platform fee and optional services
- Highlighted reasonable prices and agility of use
- Offering doubling of capacity and resiliency
- Model divided into
 - Small connectors =platform fee of \$225k
 - Large connectors =platform fee remains at \$400k
 - Rob indicated that Internet2 will be talking with the connectors at this time to assess whether they will be large or small.
- The SMS recommended that Internet2 allow for the possibility of a small annual increase (on the order of 3%) that would consistent with Internet2 increases in costs

James encouraged pag members to read the document and provide input.

Internet2 will be looking for connectors to make a pledge as to their intentions to remain large or small before year end.

IRR registration and google routing changes [Paul Howell]

- Paul reviewed Google's requirement regarding peering with them; that is, any prefixes must be in an IRR- if not, any traffic will be returned via paid commodity network. Paul indicated that the latest information from Google is that they intend to require IRR registration by mid to late fall. Internet2 has been working to educate members about how to document their prefixes. (Hurricane Electric already required IRR registration.) It was noted, that having the prefixes documented will actually aid networks' automation efforts which, in turn, helps Internet2 in our automation effort.) Paul pointed out that accomplishing this is a fundamental component of MANRS.

3. Next meeting – October 10, 2019, 11am-12pm ET.

NAOP -Platform Fee Proposal September 2019

In the Spring of 2019, the Network Architecture, Operations and Policy Program Advisory Group (NAOPpag) Next Generation Infrastructure (NGI) Service Model Subcommittee (SMS) shared a proposal with the connector/network member principals for a new Service Model that could be implemented with the Next Generation Infrastructure. The next step was for Internet2 staff to talk with members of the community about the proposal and gather their feedback. The SMS agreed to reconvene in September, 2019 to review that feedback and approve a plan of action going forward.

Internet2 staff held a webinar with the connector/network member principals in late April. Since then, Internet2 staff has had calls and discussed the Platform Fee model with several individual Regional R&E network members. In July, Internet2 presented the Platform Fee Model at the Community Architects' Workshop, a gathering of over 100 network engineers and architects from regional networks and campuses.

Rob Vietzke presented the Platform Fee Model and potential fees to the Internet2 Board of Trustees in July, 2019. The Board asked that Internet2 staff proceed with discussions with the SMS and the community focusing on the new value members will receive to support research, cloud, academics, and enterprise from this new network under the proposed Platform Fee model.

Summary of the proposal the SMS presented to the community in Spring 2019

The Network Architecture, Operations and Policy Program Advisory Group (NAOPpag) Next Generation Infrastructure (NGI) Service Model Subcommittee (SMS) is encouraging Internet2 to offer a Service Model that will bring value to the Connectors and their members by utilizing the wide range of new features and capabilities the NGI Platform will provide. There are two tracks to a new fee model: Updating the value and fees of the current "port" offerings, and introducing new services including new low-cost options for peering at major peering points. For the "port" fee updates, the SMS has recommended the development of a Platform Fee that would allow Connectors to benefit from a wide range of services, while investing at a sustainable level.

Under the Platform Fee model it is proposed that Connectors would have access to all of the Layer 2 and Layer 3 services shared by the community, such as:

- *Access to the world-class Research and Education (R&E) network to support data intensive research and the academic enterprise, including international access, updated measurement, telemetry, security and data movement capabilities (examples include NRP, OSG, ERP, etc.).*
- *Complete R&E routing table and cloud connectivity to support member-to-member and member-to-cloud engagement. (Required)*
- *Shared access to Cloud Connect Services available via the Cloud Connect Portal to major cloud IaaS providers.¹*
- *Shared access to the Internet2 Peer Exchange [I2PX] (formerly TR-CPS).*
- *Layer 2 connections across the platform.*
- *Additional capacity available as new solutions become viable (e.g., when 400G is available).*

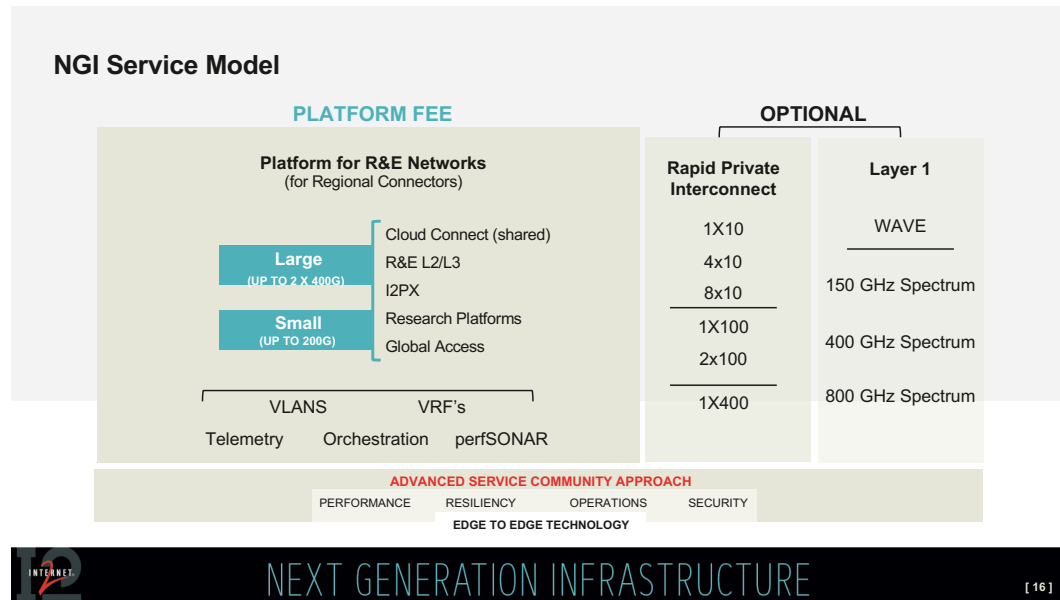
The Platform Fee model contemplates providing Connectors with flexible capacity options. The final determination will need to be informed by the selections for optical and programmable packet solutions. It may be possible to offer Connectors the ability to take advantage of capacity increases as the technology and its costs allow. For example, a Connector may be offered double their current capacity when the system is available in 2020. But that same system may be able to support 400G channels for a similar cost just a few years later. A model is envisioned where a Connector could take advantage of that capacity growth by paying only the increase in equipment costs (e.g., a one-time non-recurring (NRC) cost to obtain the 400G-capable equipment).

Community members have said that they would like such a model, taking advantage of the capabilities of the network, without incurring large fee increases to do so. In the current model Internet2 is able to grow its revenue to sustain the network when a Connector’s traffic grows and more capacity is procured. Absent that organic growth, Internet2 will need to consider ways to obtain the revenue that is necessary to sustain the network.

Therefore, the SMS is recommending that a small annual percentage increase be implemented.¹ This will provide Internet2 the revenue growth needed to augment backbone capacity, keep up with normal annual incremental operating costs, and help build reserves.²

Services Included with Platform Fee Model

The graphic below has been created and shared with community members to help depict the services that would be included with the Platform Fee, and the optional services that will be made available as part of NGI.



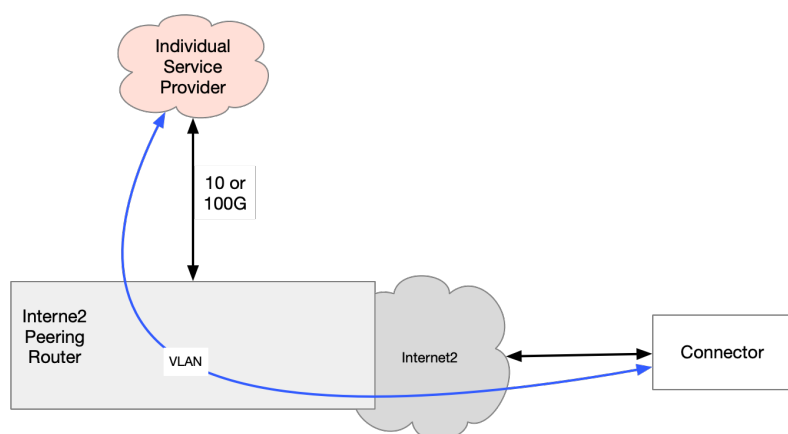
¹ In the current market environment we are considering a 3% increase per year.

² Internet2 has several contracts that require annual increases. We also spend about 25% of our annual network operating expenses on services from the Regionals, some of which also have annual fee increases.

In addition to the all the services offered as part of the Platform Fee, we are proposing two optional services, Rapid Private Interconnect (RPI), and Layer 1 Optical Services. Following is a brief description of each service. A detailed Service Description for each service is provided in Attachment A (RPI) and Attachment B (Layer 1 Optical). Each of these optional services will be offered at marginal costs.

Rapid Private Interconnect (RPI)

RPI provides a low-cost option to regional networks that they can utilize for themselves or extend to one of their members. RPI allows dedicated, private, regional-controlled access to any vendor-provided services offered at one or more of the peering exchanges. This could serve as a primary connection or redundant connection to the service.



Layer 1 Optical Service: Wave Service

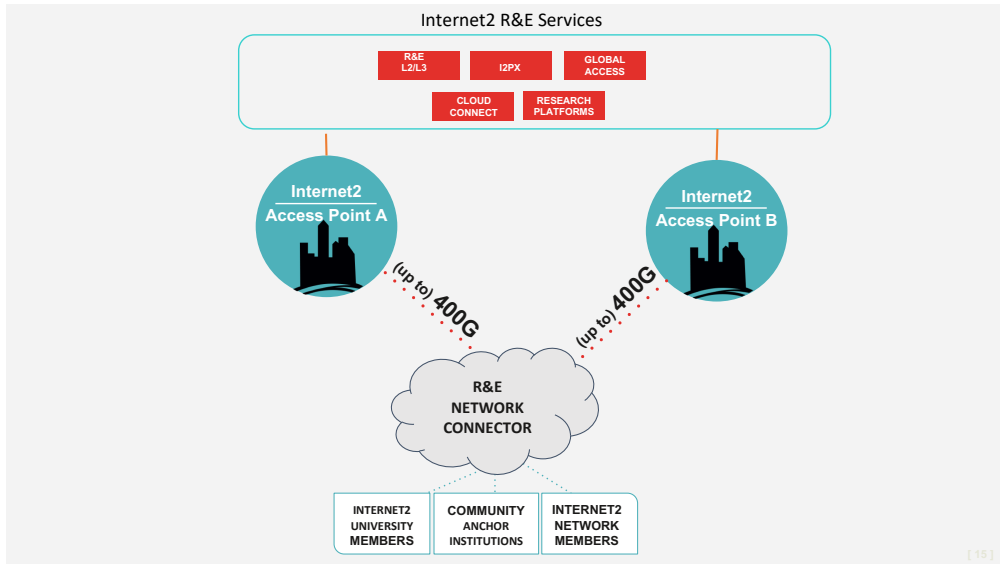
The Internet2 Wave Service provides members with the ability to provision dedicated waves on the Internet2 Network. Currently waves can be configured with capacities between 10 and 100 Gigabit. The Wave Service includes a complete solution using Internet2-supplied equipment, with maintenance and support included.

Layer 1 Optical Service: Spectrum Service

The Internet2 Spectrum Service gives participants the ability to provision dedicated waves on the Internet2 Optical System. Currently 50 Ghz channels are available in the C-Band, but flexible grid channels (or blocks of spectrum for multiple channels) will be introduced beginning in 2020. Organizations wishing to procure spectrum services supply the equipment to light the spectrum and Internet2 will install the equipment. The equipment can be either the same used by Internet2 to light the Optical System or equipment from other manufacturers with prior approval.

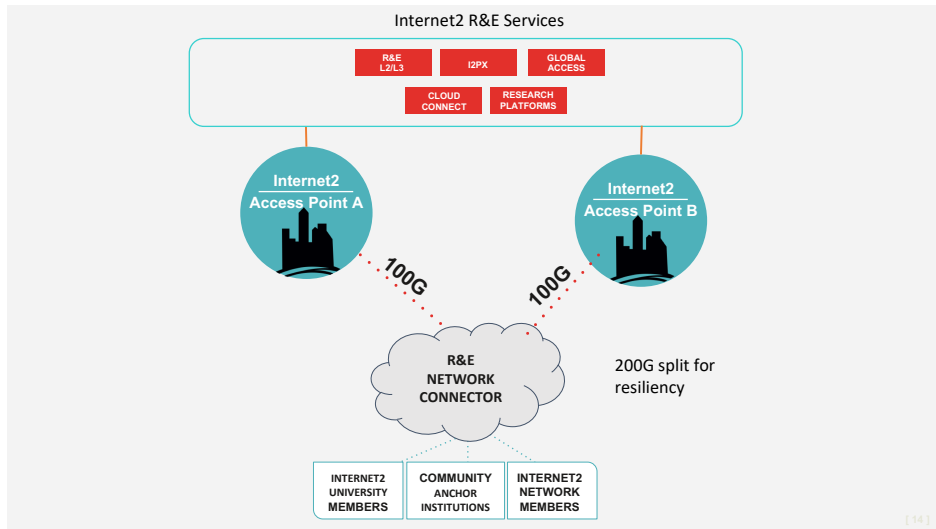
NGI Service Model Proposals

We believe we can deliver increased capacity to Large Connectors (currently with dual connections), and increase the resiliency of Small Connectors (currently with single connections) for about the same fees that the connectors pay today. With the proposed Platform Fee Model, Connectors will be able to choose Large or Small configurations. The Large configuration will provide them access to the NGI network at two locations, with up to 400G of capacity available at each location.



Large Connector Model with resilient connectivity

Those who choose a Small configuration, would have access to the NGI network in two locations with 100G of capacity at each location (this model is designed to provide the resiliency that many Connectors stated was an important goal for them with NGI).



Small Connector Model with resilient connectivity

Currently, Internet2 has 34 Connectors.

18 Connectors (Single Connectors) have a single 100G port today for which they pay \$200K per year.
 16 Connectors (Dual Connectors) have dual 100G ports today for which they pay \$400K per year.

We have analyzed several scenarios for setting fees for the Large and Small Platform models. We are proposing that we use the following scenario to provide Connectors the potential fee for the Large and Small Platform. We would ask Connectors to confirm their platform level, then apply that information to the business modeling.

Proposed Model	Quantity	Annual Platform Fee
Small Connectors	18	\$225K
Large Connectors	16	\$400K

Appendix A

Rapid Private Interconnect (RPI) – private connectivity at peering sites nationwide.

RPI Service Description

Rapid Private Interconnect (RPI) allows Internet2 connectors to present themselves for private peering at selected national peering locations. Connectors negotiate their own interconnection and peering without additional investment in new backhaul, colocation, equipment or management overhead. The service will utilize the increased bandwidth and resiliency available in the Next Generation Infrastructure (NGI).

Internet2 currently operates in seven major peering points, geographically dispersed throughout the country (see Table A). These locations have been used to support much of the external connectivity to provide the Internet2 I2PX peering service, NET+ providers, and international R&E networks. As part of the NGI program, Internet2 is augmenting these peering points to offer the Rapid Private Interconnect (RPI) service for Connectors.

What is the offering?

RPI provides a low-cost option to regional networks that they can utilize for themselves or extend to a one of their members. RPI allows dedicated, private, regional-controlled access to any vendor-provided services offered at one or more of the peering exchanges. This could serve as a primary connection or redundant connection to the service.

Working with regional R&E networks, Internet2 can add a significant level of resilience and provide “out of region” connections for members of regional networks. The following are some examples:

- **Northeast** – regional members in New England, who provide local resources in New York City could use RPI to provide connections in Ashburn, VA or Chicago for their members.
- **North Central** - regional members who provide local resources in Chicago could use RPI to provide connections in Ashburn, VA or Dallas for their members.
- **West Coast** - regional members who provide local resources in Sunnyvale could use RPI to provide connections in Dallas or Seattle.
- **South West** - regional members who provide local resources in Dallas could use RPI to provide connections in Chicago or Sunnyvale

RPI solutions provided to date

Below are some examples of several Proof of Concept projects Internet2 has been developing with community members:

- A regional network in **New England** received a request from a university member to provide direct connectivity to the Oracle Rapid Connect service in **Ashburn, VA** to support the migration of their ERP system from a campus-based system to Oracle's cloud. Using a 10G RPI vlan over Internet2 the regional connected to Ashburn, where a cross connect was installed from the Internet2 router to Oracle's Rapid Connect port.
- A **mid-Atlantic regional** network had requests from colleges initiating E-Sports programs to provide direct access to a gaming internet exchange in **Ashburn, VA**. The regional is using an RPI vlan to reach Ashburn, where they connect to the E-Sport provider.
- An Ivy League university had made significant investments in dedicated network connections to AWS. The university now plans to use RPI to provide transport from their campus to **Ashburn** via Internet2, then will cross connect to an AWS Direct Connect port. The university required resiliency to reach multiple AWS Direct Connect ports and is using RPI to provide redundancy and leverage their Internet2 capacity at a significantly reduced cost.

How will this be configured?

The RPI service is available today using a manual configuration, we have future plans to make the configuration available using the OESS-based *Cloud Connect Portal*. As the service matures over the next two years, we will continue to add services to the OESS *Cloud Connect Portal* to eventually accommodate a full set of routing, security, and other parameters that connectors may wish to manage at the peering points. The RPI service is being offered today utilizing 10G ports already in place. Options for RPI 100G ports will be available as needed.

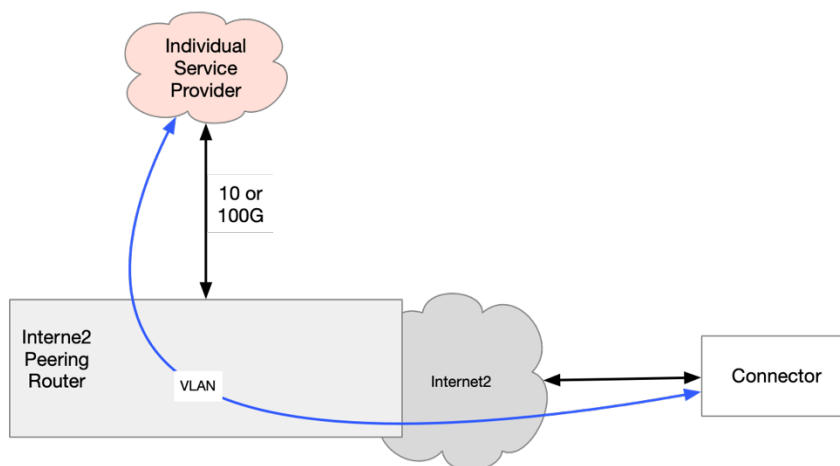


Table A

Locations:

- Equinix, Ashburn, VA
- Equinix, Cermak, Chicago, IL
- Wilcon, 1 Wilshire, Los Angeles, CA
- Equinix SV5, Silicon Valley, CA
- Telex, 111 8th, New York, New York
- Equinix, Datamart, Dallas
- SIX, Westin Building, Seattle

Technical Specifications:

- 10G single-mode LR (10km) ports
- Up to the capacity of the combined RPI ports can be transited on the Internet2 backbone to a connector's 100G/400G ports

Layer 2

- Layer 2 vlans through **Cloud Connect** Portal
- Multi-port LAG *
- Reserved Bandwidth

Layer 3

- VRF's through **Cloud Connect** portal (Internet2 AS)
- VRF's using Regional or campus AS *
- Multi-port LAG *
- BGP Prefix Lists *
- Anti-spoofing filters *
- Netflow Data *
- Routing Policy *
- Reserved Bandwidth [does portal do this?]

* Internet2 will manually provision this capability at the start. *Cloud Connect Portal* configuration is planned in the future.

Note: These RPI ports are not intended to connect member traffic or to provide transit between Internet2's other services in replacement of connector ports.

Appendix B

Internet2 Advanced Layer 1 Services

Service Definition: Wave Service

Organizations wishing to procure the Wave Service are provided with a complete wave service using Internet2-supplied equipment in the waves, with maintenance and support included.

Requirements for Participation in this Service

To participate in this service an organization must be an Internet2 member and must be able to cross-connect to the waves at the locations specified for the origin or termination of the waves.

Service Characteristics

Wavelengths are available as 10 and 100 Gbps. 10Gbps waves are carried as subchannels on 40G or 100G transponders. An optical client interface is provided to the customer with 10KM single mode fiber interface. Waves are provisioned on a single point-to-point fiber path. Optical protection redundancy and longer reach client optics are available at a premium fee. Encrypted waves are available at a premium fee. The network supporting these waves is provisioned from carrier-class facilities and benefits from a separate management network and security program protecting the infrastructure.

New features that are planned for as part of NGI include:

- Custom telemetry streamed to the member for each provisioned wave and its path.
- Flexible grid support for channel and symbol rates other than 35 GBaud at 50 Ghz.
- Support for equipment from third parties other than Internet2's primary vendor.

Service Definition: Spectrum Service

The Internet2 Spectrum Service provides members with the ability to use their own optical transponders on the Internet2 Open Line System to provision their own wavelengths on the Internet2 Network.

Currently spectrum can be configured with capacities between 10 and 100 Gigabit. Beginning in 2020, spectrum will be available in flexible increments to allow various baud rate transponders to be used.

There are more than 50 locations on the 16,000-mile fiber footprint where spectrum can originate or terminate.

Organizations wishing to procure spectrum services will provide a single fiber pair interconnection from their own optical equipment to the Internet2 line system. If the user wishes multiple ports on the Internet2 ROADM to break out their spectrum, a small additional fee may apply. The equipment can be either the same as used by Internet2 to light the Optical System or equipment from other manufacturers with prior approval (to allow modelling in system tools).

Requirements for Participation in this Service

To participate in this service an organization must be an Internet2 member and must be able to cross-connect to the spectrum at the locations specified for the origin or termination of the spectrum. The organization purchasing the service must procure equipment to generate the wavelengths, coordinate with Internet2 NOC on channel management procedures and have plans for maintenance of its equipment.

Service Characteristics

Network Media Channels (NMC; Optical Spectrum allocations) are configured on an optical path on the Internet2 network. The member's optical equipment is connected to one or more NMC connection points on the Internet2 optical network. Internet2 manages the NMC between the endpoints of the service and the member manages the optronics that put optical light into the NMC. Outages may result in the case of a fiber cut or equipment failure. Premium Options include asking Internet2 to manage the customer-purchased equipment and/or provisioning of resilient paths. The network supporting these waves is provisioned from carrier-class facilities and benefits from a separate management network and security program protecting the infrastructure.

New features that we plan to introduce as part of NGI include:

- Custom telemetry streamed to the member for each provisioned wave and its path.
- Flexible grid support for channel and symbol rates other than 35 GBaud at 50 Ghz.
- Support for equipment from third parties other than Internet2's primary vendor.