June 2024
Large Platform Flex Committee Report

Executive Summary

Internet2 created a committee of members to review and advise on a new service offering from Internet2 called Large Platform Flex. The committee met several times and provided Internet2 staff with ideas on how to improve the services, and provided several concepts on how the service might be used by community members. This report provides the results of the committee's work, outlines the Flex service and the improvements suggested by the committee. We also offer, herein, some use cases for the Flex service that are already in place, or being contemplated. Lastly we provide some background on how the five-year long NGI effort to upgrade and automate the Internet2 infrastructure has led us to the ability to offer the Flex service.

Initially, the Flex model was intended for Large Platform subscribers only, but the Committee suggested that there are many cases where Flex interfaces would be important to Small Platform subscribers as well, so Internet2 took the advice of the Committee and is now offering Flex interfaces to all Platform subscribers. Thus the name of the service has changed from Large Platform Flex to Platform Flex.

The Platform Flex service is available to Internet2 Network Connectors (those with physical connectivity to the Internet2 network). The Platform Flex model is designed to allow Connectors to offer their members more connectivity options and to leverage the automation tools that have been developed as part of NGI. Connectors will be able to utilize the Internet2 Insight Console to manage, troubleshoot and monitor any of their Flex connections, and at their discretion, may decide to provide access to the Insight Console directly to their members for self-provisioning and self-management. Platform Flex interfaces will only be made available to Internet2 members or Research and Education Network (REN) members through the Connectors.

Please join us in thanking the members of the Large Platform Flex Committee for their contributions and for improving the service with their knowledge and experience.

Background on Large and Small Platforms

In 2021, Internet2 introduced the Platform model for Internet2 Network Connectors. The **Platform Model** offers two options:

- **Large Platform** – up to 800G total capacity across two locations (through 4x100G or single 400G interfaces) accessing Internet2 R&E, Internet2 Peer Exchange (I2PX), Internet2 Cloud Connect (I2CC), Global Access, and Research Platforms
- **Small Platform** – up to 2-100G connections accessing Internet2 R&E, I2PX, I2CC, Global Access, and Research Platforms

Introduction of Platform Flex

In 2023, Internet2 introduced **Platform Flex**, allowing more interfaces in each city where a Network Connector utilized Platform ports (Home City) to support Connectors with better resiliency, additional services and utilization of Internet2 network capacity. Platform Flex introduces the ability for Network Connectors to utilize additional interfaces at no additional fee in each of the Home Cities in which they connect to Internet2. (For more background on what led to Internet2 being able to offer Platform Flex please see Appendix B.)

The interfaces could either be additional 100G interfaces for those utilizing 100G interfaces, or a second 400G interface for those utilizing a single 400G interface. All interfaces **utilize the existing 100G or 400G of Internet2 backbone access capacity** currently allocated in a Home City through a Platform subscription.

Connectors may purchase additional interfaces, if they wish to use more than the allotted additional interfaces in one of their current Home Cities, for $8,000 each per year. The interfaces can be used for access to the Platform capacity in each Home City or to connect to local resources. Local resources might include connections to an internet exchange point or caching servers located locally with the Platform connection. If a Connector needs connections in a city different from their current Home Cities, the Connector may procure additional backbone capacity in any city on the Internet2 Network for an additional fee.

All Flex ports are able to be managed by the Virtual Networks feature of the Internet2 Insight Console.

Large Platform Flex Committee*[1] Charge

In 2024, the Large Platform Flex Committee (LPFlex Committee) was formed with the following community members:
Jon Ellis, FLR
Chris Griffin, FLR
Scotty Strachan, NevadaNet
Tracy Smith, University of Illinois, Urbana-Champaign
Lance Taylor, FLR

The Large Platform Flex Committee was charged with working with Internet2 staff to review uses of the Large Platform Flex and provide additional input to the service, including describing use cases for the Internet2 Large Platform Connectors as they begin to expand the usage of their Flex connections.

Work of the Committee
In one of the initial meetings, the members of the LPFlex Committee spent time discussing use cases for the Flex model across the community. These discussions illuminated use cases where Small Platform subscribers could benefit from the Flex model as well. As stated by committee member, **Scott Strachan**;
"Some applications need features but not necessarily capacity… allowing small platform users access to Flex will enable them to use it creatively." Internet2 is heeding this advice and has agreed to offer the Flex option to both Large and Small Platform subscribers.

Committee member **Tracy Smith** cited how Platform Flex might provide advantages to research organizations, stating; "***campuses and entities requiring secure enclave environments that will meet forthcoming Cybersecurity Maturity Model Certification (CMMC) 2.0 requirements for regulated research, may find Flex options useful for segregating traffic to secure cloud service offerings."***

Below are existing and potential applications for use of Platform Flex that were identified by Internet2 staff and committee members:

- Redundancy for those using 400G

[1] The Platform Flex model was initially planned to be offered to Large Platform subscribers only. However, through the work of the Large Platform Flex committee it was suggested that Platform Flex be offered to all connectors regardless of whether they subscribed to a Large Platform or a Small Platform. That suggestion has been adopted, so the moniker of Large Platform Flex has been changed to Platform Flex.

    - UETN and MAX have expressed an interest in implementing 400G redundancy
- Dedicated 100G interfaces for specific purposes
    - Dedicated to I2PX or R&E, or Research or cloud
    - Dedicated to Local Service Providers
    - MACSEC implementations
- Research Support
    - Flex interfaces dedicated to specific research efforts
    - Science DMZ implementations
    - Secure Enclave
    - CC* applications
    - Interfaces dedicated to specific organizations (e.g. HPC Center)
        - Dedicated Research Interfaces
        - Research-based partnerships could have a different connection, tied to a research project, supporting a multi-organization project team tasked with design implementation across the project[2]

- Virtual Network and Cloud Connectivity
  RENs could allocate Flex ports to institutions who request them for:
    - Allocating an interface to cloud connectivity
    - Facilitating the use of Insight Console Virtual Networks features
    - Utilizing Insight Console capabilities to manage their cloud connectivity
    - Utilizing Insight Console capabilities to manage their research collaboration with other institutions
    - Utilizing upcoming Cloud Router capabilities to manage their virtual networks directly
    - Obtaining a redundant path that does not include REN equipment

- Value Proposition Considerations for RENs
    - Flex interfaces will be made available only through Connectors
        - It is the investment in the Platform made by the Connector that makes the Flex options available
    - RENs have the option to charge a fee for the use of Flex interfaces
        - Helps offset the cost of the Platform Fees
        - Helps cover the cost of any additional resources the REN utilizes to deliver the Flex capability
    - RENs may be able to charge a fee for additional local ports or circuits to connect institutions to the Flex interfaces

[2] Multi-organization project teams could take advantage of Flex interfaces in combination with the I2 Insight Console virtual networks and Cloud Router capability to establish and manage connectivity between the sites housing the teams.

    - RENs may utilize Flex interfaces to connect to local service providers, (e.g., content providers, local IX, local SaaS service providers), delivering additional valuable services without additional equipment costs

- Flex utilization guidelines for RENS
    - Flex ports are ingress ports in Home Cities where a Connector is located;
        - If a Connector wants an egress port to get off an Internet2 router to connect to another regional in another city they would use RPI
        - Flex Ports are only available to a Connector in their Home Cities
        - Each Connector can use their flex ports, in their home cities, as they see fit, including using one of their Flex ports to provide egress to another Connector with Internet2 approval

- Communicating that the key to this offering is Flexibility
    - The Committee agreed that it was important to communicate to the Connectors and institutions that they should feel comfortable bringing forth any (legal) use of Internet2 interfaces, even if they are not listed as approved

Appendix A

Use Cases

NEREN Use Case for Flex to Deliver Layer2 Local Exchange Operations

The NEREN use case shows a model where Internet2's newest Connector is taking full advantage of the Platform Flex and I2 Insight Console capabilities to fully leverage the Internet2 NGI routing platform resources. Read on about this bold initiative to utilize the flexibility and automation capabilities of NGI to operate a Connector without the cost of owning one's own equipment.

Internet2 will offer a virtual private switch capability to allow NEREN to run its own "private exchanges" on the Internet2 switches in Albany and Boston. NEREN would like to avoid running its own exchange switches, operations, etc and instead use Flex ports on the Internet2 switches. NEREN would like to be able to do anything on these ports that it could do on their own Layer 2 switch. NEREN does not require layer 3 exchange. NEREN members could run a route-server or router on a stick on one of their own devices if need be.

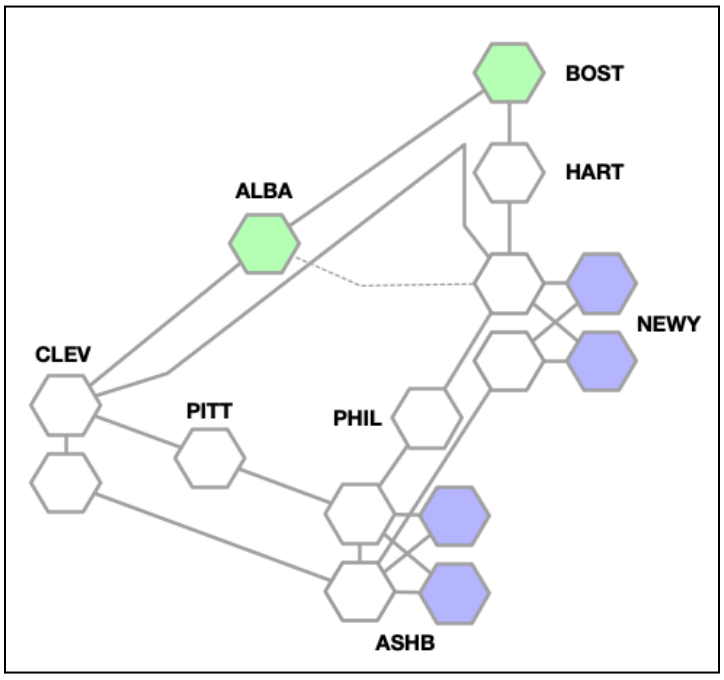

The image above shows router sites in the NEREN footprint (green) as well as likely backup paths and cloud interconnect locations (purple).

Services delivered to NEREN:

NEREN individual state members will be able to create VLANs to Internet2 services from these Flex ports

- VLAN to each NEREN member for Internet2 services (similar to Big Ten Academic Alliance Setup in Chicago). Likely a set of Internet2 R&E, I2PX, DDoS, cloud connect, etc.

Port to Port and Port to multiport VLANs at Layer 2

- NEREN member to NEREN member VLANs for private in-region purposes
- NEREN shared VLANs for special purposes, perhaps, for example, the Massachusetts Green High-Performance Computing Center (MGHPCC). A VLAN to which all NEREN members might connect. Routing would be hosted by one of the NEREN members
    - Example 1: An example of one-to-many might be a trunk to BOSTON-IX with multiple VLANs coming through the port and then fanned out to NEREN members.
    - Example 2: Another example might be NEREN members purchasing a 100G port with a commodity service provider, and bringing in a VLAN per member to that port and fanning them out to other NEREN member ports.
- NEREN will be able to utilize the I2 Insight Console to create VLANs along the short path from Albany to Boston using the NEREN backbone and also the long way around from Albany to Boston using Internet2's 400G backbone.

NoX Managed Infrastructure Pilot

The Northern Crossroads (NoX), in discussions with Internet2, would like to collaborate and see how far NoX can leverage the Internet2 infrastructure.

After learning about the Internet2 Platform Flex (PF), NOX engaged with Internet2 to pilot Flex ports with some of their large member universities. The institutions are using NOX Flex ports to connect directly to the Internet2 NGI packet network to provide cloud connectivity via the Insight Console Virtual Network service. NOX also intends to explore leveraging some of the early features of the Internet2 Cloud Router Pilot to facilitate local peering relationships.

Internet2 and NoX intend to leverage the Platform Flex (PF) with Insight Console to collaborate on a pilot to incrementally develop additional capabilities supporting the NoX membership. This is a great opportunity to work with a regional that is interested in using PF and connecting a select number of schools to utilize the Insight Console. This will ideally lead to an opportunity to spotlight the work and use cases with other Connectors, regional networks and Internet2 members.

There is great potential in success for Internet2 and the RENs should this pilot to share infrastructure services prove successful. Such a service could be offered to like-minded RENs upon request. Also such a service would make it easier for institutions to fully utilize the advanced network automation capabilities of I2 Insight Console.

Out of Region Redundant Connectivity Requirements

Internet2 and a Connector were approached by an institution to seek a solution for a peering connection to a provider that was located outside of the home region of the institution. A main criteria was complete redundancy, which meant not utilizing the Connector's equipment for the redundant path.

Internet2 is working with the Connector to utilize a Platform Flex port to provide the institution with a redundant connection into an interface on an Internet2 router, thereby avoiding the use of the Connector's equipment. The institution will be able to leverage Internet2 reach beyond the region and utilize a port on Internet2 equipment to deliver the required redundancy, while still leveraging the investment they had made in their membership with the regional Connector, because this service was only available to the institution through their Connector.

Appendix B

The Flex Offering is the Result of Delivering on the Promise of NGI

In February 2018, Internet2 presented the following statements at the Connector and Research and Education Network Member Principals' meeting regarding the goals for the NGI effort. (The items in bold are updates to the original message that identify how Internet2 has achieved the stated goal.)

The Next Generation Infrastructure Program (NGi) is focusing on providing software driven solutions for data-intensive research and enhanced cloud access. The NGI project incorporates updated services and increases value with a lower operating cost for the supporting technology of the Internet2 infrastructure portfolio (and relationships in the larger ecosystem).

- *Includes the services and service models through which the community adopts Internet2 infrastructure services*
    - ***We delivered the Platform model, RPI, and Cloud Connect services.***
- *Includes new features, primarily driven by software, automation, and systems virtualization to allow the infrastructure to be more readily integrated into the broader campus, regional, and cloud environment around us.*
    - ***We delivered I2 Insight Console services including Looking Glass, Community, and Virtual Networks allowing broader integration.***
    - ***Internet2 has deployed the NSO network automation orchestrator and now operates its production network via automation. Engineers no longer log into any production routers.***
- *Includes a number of infrastructure upgrade projects that reduce the operating expenses of the network infrastructure.*
    - ***We delivered an automated optical, packet, and software platform that significantly reduced our operating costs (power was reduced 70%) and enhanced our network management capabilities.***

In 2024 we added the Flex option to Platform subscriptions to help the community utilize the capabilities that NGI has delivered. Below we identify the foundational improvements we have implemented with NGI that have led to the ability to offer Flex to the community. These represent the building blocks that led to improved services and value the community can now leverage in new and exciting ways in the support of data intensive research, enhanced cloud access, routing integrity, and systems virtualization.

1) A big part of NGI was an automated infrastructure that was much more affordable to operate
    a) Lower power
    b) Smaller footprint

    c) Right-sized packet platform (including “Licenses for Use” for all existing interfaces)
    d) Automation
2) Implemented an NSO model which automates the management of the packet network, ensuring a uniform implementation of packet network policy and easing the network management burden. With NSO we could offer and manage more connections to the NGI packet network without significantly increasing the workload.
3) Prior to NGI, network providers like Internet2 would deploy chassis-based routing equipment. Providers would populate the chassis with the appropriate amount of interfaces needed, then would procure and slot more cards in the chassis as capacity demanded. Each additional card came with additional costs and these chassis required more power even when not fully populated. (Increased use meant increased capital and operating costs.)
4) The NGI packet network used fixed-format devices, a model that is very much in pace with what hyperscale providers deploy (the target market for these devices). So the NGI platform has a fixed number of interfaces that can be deployed without any additional hardware costs.
5) Many router vendors offer fixed-format packet gear and offer a ‘license for use’ for each port that is to be utilized. So one only pays for the interface if it is in use. But again, growth in the number of interfaces in use increases the costs (just as in the pre-NGI chassis model - #3 above).
6) Internet2 has always worked on a model of providing abundance. Our service models allow for bursting, and we overprovision our backbone to meet elastic demand. So when selecting the current NGI platform, we ensured that all interfaces would be available at no additional charge when they were utilized. This model means no increase in the cost to Internet2 if more interfaces are used, and allows Internet2 to encourage use of more interfaces through the Flex model.
7) Internet2 has been working with the community to develop the Insight Console, the landing page for all I2 Network Services and the place to manage, troubleshoot and visualize the I2 network. Insight Console is being enhanced constantly, and its current features support troubleshooting, creating community, routing integrity, and managing virtual networks. Flex is leveraging the functionality of I2 Insight Console to allow members of the community to configure and manage virtual networks. These capabilities allow the infrastructure to be more readily integrated into the broader campus, regional and cloud environment around us.