We are looking forward to seeing you next week for our annual gathering to set direction for the upcoming year. The focus of the meeting will be to provide an update on the progress being made on the Next Generation Infrastructure (NGI) effort, and to seek your input on important decisions about the services we plan to offer with the new capabilities NGI may provide. Lastly, we will use the experiences of one campus' migration to the cloud to discuss how the community might prepare to support these services across the ecosystem through automation and coordination.

We are pleased that all of the members of the NAOP NGI Subcommittee on Service Models will be in attendance. Our goal is to have these community members hear directly from our Connectors and Network Members about their views on the value of proposed future services and the service models we may employ to fund them. The Subcommittee is chaired by Dee Childs, who along with Jim Stewart will also serve as our co-chairs for the Tempe meeting.

In preparation for the meeting, we ask that you read the background information provided before the meeting next week. We have made suggestions as to which is most important for you to preview.

As stated earlier the meeting is made up of three parts:

Tuesday morning: ***Internet2 Next Generation Infrastructure Project: Where We Are***
During this portion of the meeting we will

- Review current NGI progress and financial assumptions (background is provided in the document called "2022" Community Infrastructure Services Capabilities & Planning Direction (Update) beginning on page 3; **this paper should be reviewed.)**
- Begin a conversation about how to introduce the new Next Generation Infrastructure (NGI) service capabilities and increase value received by the connectors. As background for these conversations, we have provided description of current network services (beginning on page 13) and created a set of potential service descriptions (beginning on page 23; p**lease read one or two of the service descriptions, perhaps one current and one potential new service.)**

Tuesday afternoon: ***Gathering Input from the Principals on Services***
We will spend the afternoon talking about how to use these service definitions to create a service model that would address new fees going forward. We will present some service model options for our discussion (see document called Potential NGI Service Models beginning on page 30**; we ask that you review this document to prepare for the discussion)**. We expect your engagement to provide input to the process to transition to one of more of these new service models.

Wednesday morning: ***Looking Forward with a Focus on Cloud, Automation and Ecosystem-wide Delivery***

Finally, Wednesday morning will provide an opportunity to look forward to how we will provide these services across the ecosystem through automation and coordination of our services. We have asked a campus leader who led the transition from a campus data center environment to an environment entirely in the cloud to provide some concrete stories about the end-to-end challenges. We will end with a discussion about how we can work together in 2019 to build support for ecosystem wide needs.

DRAFT for discussion with Principals

2019 Tempe Connector/Network Member Principals' Meeting
January 15-16, 2019

Agenda
January 15, 2019

8:30-9:00	Breakfast
9:00-9:45	**Welcome** (Meeting Chairs: Dee Childs (TAMU), Jim Stewart (UETN)) **Host City Welcome** (Lev Gonick, Chief Information Officer, Arizona State University) Introductions **New Principals**: Matt Riley (Oregon Gigapop), Dan Schmiedt (C-Light) **Internet2 Board Regional Board of Trustees Members**: Louis Fox (CENIC), Marla Meehl (FRGP) **NAOPpag NGI Service Model Group**: Dee Childs (TAMU; Chair), Cort Buffington (KanREN), Steve Kankus (NYSERNet), Matt Riley (U of Oregon), Paul Schopis (OARnet), Marc Wallman (U of North Dakota) **Internet2 Welcome** (Howard Pfeffer)

Tuesday morning focus: Internet2 Next Generation Infrastructure Project: Where We Are

9:45-11:15	**Next Generation Infrastructure Overview, Progress and Plans** (Rob Vietzke) Update of NGI project and community progress in 2018 Update on Revenues and Expense projections / assumptions for service models
11:15-noon	**Internet2 Service Descriptions** (Rob Vietzke, George Loftus) Current services How NGI related features, architecture changes and other ideas might map to services
Noon-1:00	Lunch

Tuesday afternoon focus: Gathering Input from the Principals on Services

1:00-1:15	Recap from the morning Service Descriptions (Jim Stewart, George Loftus)
1:15-2:00	Break-out session 1
2:00-2:20	Feedback from Break-out session 1
2:20-2:50	Refreshment Break
2:50-3:40	Service Models Discussion (Dee Childs, George Loftus)
3:40-4:20	Break-out session 2
4:20-4:40	Feedback from Break-out session 2
4:40-5:00	Wrap up the day and a look ahead to Wednesday
5:00-6:00	Reception, Tempe Mission Palms
6:30-9:00	Dinner, Caffe Boa

January 16, 2019

8:30-9:00	Breakfast

Wednesday morning focus: Looking Forward with a Focus on Cloud, Automation and Ecosystem-wide Delivery

9:00-10:00	Presentation: "Cloud First: a hands-on look at the impact of cloud moves on higher education" (David Seidel, VP for Information Technology and CIO, Miami University)
10:00-10:30	Break
10:30-11:30	Discussion
11:30-noon	Wrap up and next steps
Noon-1:00	Box Lunch (Seating in the courtyard available)

INTERNET2 NETWORK SERVICES

DRAFT Service Catalog for discussion with Principals

CURRENT INTERNET2 SERVICES

POTENTIAL NEW NEXT GENERATION INFRASTRUCTURE SERVICES

DRAFT for discussion with Principals

"2022" Community Infrastructure Services Capabilities & Planning Direction (Update)

Preface:

January 2019
The following is an updated version of this document that was first presented to the community at Global Summit 2018, in May 2018. We have included any updates on progress made since Global Summit 2018.

The Research and Education networking community is being asked to address several issues facing all of us with the implementation of the Next Generation Infrastructure. These include:

- Expanding capacity offered to regionals, including removal of TR-CPS utilization caps, introduction of private cloud access, and expansion of the capacity for research;
- Emerging research-support opportunities like implementation of infrastructure and programs for a national research platform;
- Introduction of automation, self-service tools, and API's, all to enable the sharing of community infrastructure and ease-of-use for users of end to end services; and
- Delivery of these features and improvements with the goal to lower total cost to operate the network for the services that it delivers.

Addressing these issues will take time. Beginning now allows us to achieve solutions to these issues by 2021/22. Building on 3 years of community discussion, our meetings at the 2018 Global Summit were a pivot point to begin the transition from discussion of requirements to articulating a vision of the services that the community seeks to deliver as its infrastructure is renewed. Our process to date has been robust: the planning process began with discussion of the community's ecosystem-wide objectives, continued through architectural papers and infrastructure sharing discussions and continues with multiple regionally-led experiments that model future needs and collaborations.

It is important to note that there are increasing economic drivers for Internet2, regional networks and campuses to meet growing capacity demands in increasingly tight financial circumstances. Collectively, portions of the community's infrastructure are reaching the tail end of their useful lives and further investment in current technologies to support expanding requirements has diminishing value. Allowing the network planning to begin, and the eventual implementation to proceed keeps us ahead of these real issues.

Since this document was presented at the 2018 Global Summit, summarizing the capabilities and planning directions, we have taken the **next steps to form technical teams, to begin market surveys, develop an optical RFP, create more detailed plans and conduct more community checkpoints along the path to implementation of new services.**

> The architecture and operation of the Internet2 national backbone should be driven primarily by its purposes, good network engineering principles, and cost-effective operation. We should build the best network we can afford." – Guiding Principles for Network Planning developed jointly among Internet2 and Regional Networks in Tempe, 2016

The Infrastructure Support Mission of Internet2:

The Internet2 network, and the collaborative networks with which it connects, exist to provide advanced capabilities in support of research and academic collaborations across the US. Our collective missions are multi-faceted and responsible for providing not only infrastructure for the academic enterprise, but also research cyberinfrastructure for various data intensive science domains and support for network research. Together, these three service areas must be supported to enable the United States academic community.

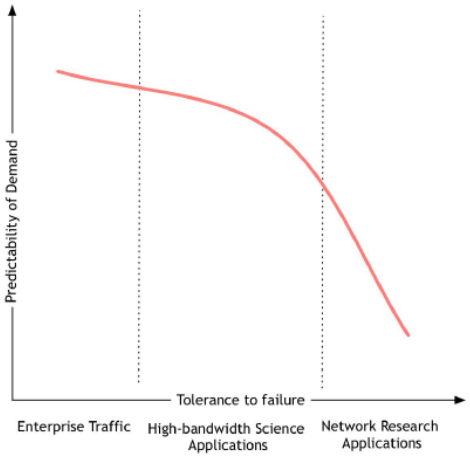

Figure 1 - "The demands of each category are very different". Ilya Baldin and Mark Johnson, RENCI - RECINNS Paperç

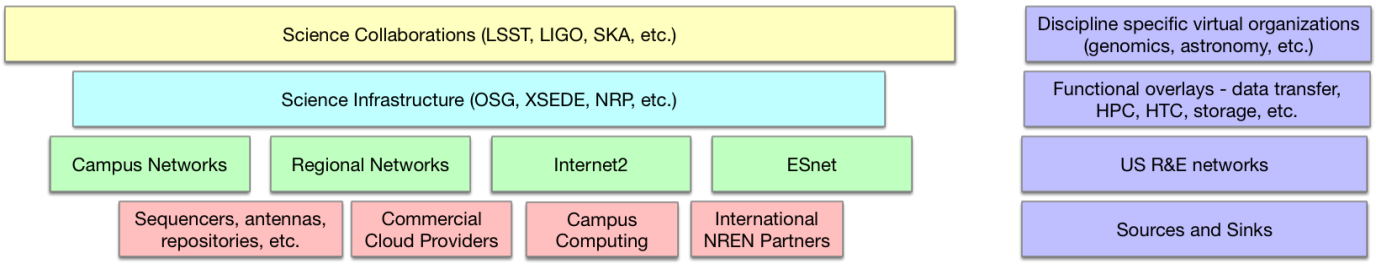

Figure 2. CI ecosystem conceptual diagram

High Level Summary of Future Requirements:

For Researchers, the network must accelerate the rate of scientific discovery by allowing researchers to get the data they need, where they need it, and when they need it. Building on successful models like the Pacific Research Platform, LHC-ONE and others, the network must provide a national facility for high-performance data movement, distributed HPC and HTC among researchers, their instruments and their collaborators. To accomplish this, the Internet2

network must operate like an uncongested freeway system with capacity for research allocated on multiple 100G+ optical lightpaths. The network must connect data generators and users of that data in a national platform optimized for performance among big data instruments, research consumers, cloud computation and virtual organizations (including organizations like NRP, OSG, ERRP, XSEDE, GENI, etc.). The key is developing an end-to-end overlay model with the regionals that can support all of these efficiently with lightpaths, VRFs, VLANs, etc.[1]

> … continuing to innovate in developing and implementing advanced networking architecture and technology required by data intensive science domains will be an important theme." - Joe Mambretti – MREN

For infrastructure providers that support research Virtual Organizations (VOs) - Funded projects and organizations such as the Open Science Grid, XSEDE, GENI, etc. can be viewed as overlays on the national R&E network infrastructure - they provide mechanisms to orchestrate resources distributed across multiple campuses in support of specific research collaborations. The network ecosystem – campuses, regionals and Internet2 - must provide end-to-end services that can integrate with these "overlay infrastructures" in order that they may provide bespoke services to support virtual science organizations such as LIGO, LHC, etc. Example service parameters could be privacy (e.g. a virtual private network), security (end-to-end encryption), optimized latency (traffic engineering), etc. driven by the specific workflow demands of the science.

For Campus IT and Administration (via Regionals): the Internet2 network must deliver three easily provisioned, consumed and supported advanced networking capabilities to regionals for consumption on the campus, that provides users, administrators and providers with features that enhance service usability, controllability and reporting:

- **Performance-optimized capabilities for researchers** as described above, including the current high performance national and global network for academic entities or some future even higher performance enhancements,
- **Secure & private connections for administrative uses** including direct cloud connections and support for remote campuses;
- **Optimized implementation for administrative**, management and financial workflow support; and access to peering and content cache services in a resilient manner that supports low cost and high-performance access to content providers.

> "The distinction between commodity and research traffic destinations will continue to blur as the use of cloud services proliferate." - From Joe Breen, U. of Utah presentation at RECINNS workshop

[1] Portions of this requirement statement adopted from various Pacific Research Platform Workshop presentations.

For regional networks, Internet2 must deliver the components that each regional requires to support its broad set of connectivity needs on an end-to-end basis, both inside and outside of its region, at the wavelength, packet and services layers. The economics of the services should be such that the regional can consume services from Internet2 at costs representing the scale of the entire US R&E community. The services should include a suite of end-to-end services and be available to the regional in a way that can be integrated with their own services or passed-through depending on the regional's preference. Planners for these services should consider emerging industry standards for interoperability and end-to-end perspective, not only for network-to-network services, but also for end to end cloud and research applications[2]. The suite of building blocks and services must at minimum support researchers, campuses and other network users like K12, Libraries and other community anchors. Internet2 and the regionals, planning together, must deliver these services in the most cost-effective manner, finding efficiencies and renewing their ability to deliver more capabilities for less cost when possible.

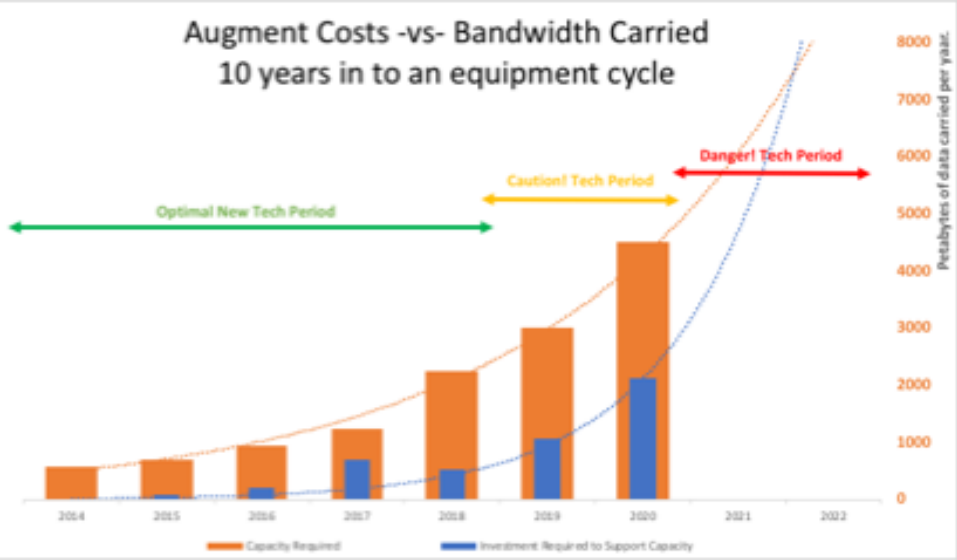

Today's R&E environment is changing rapidly. Traffic demand is growing, and fulfilling that demand with current technology platforms is becoming increasingly difficult to accomplish affordably. – RECINNS Summary Paper

[2] The Metro Ethernet Forum's interoperability standards for federated provisioning may be a good example. This is a standards-based innovation that allows cross-organization provisioning and end-to-end support.

Ecosystem Future Service Framework Declaration (draft)

As its infrastructure and supporting services are renewed, Internet2 will endeavor to offer an enhanced and simplified suite of services, with embedded security features, for researchers, academic collaborators and regional networks. These services will be factored into network building blocks that can be used independently or together in comprehensive network services. The service frameworks would allow for automated ecosystem-wide service delivery through secure software-driven infrastructure[3]. The services would offer secure integrated interconnection of big data resources among community, cloud, commercial and virtual collaborations. The service definitions are intended to be interoperable with commercial cloud and community-based services; and, interchangeable with service definitions that regional partners may declare as their own services in support of joint members and other entities they serve.

The services would include the following core components, which could be configured both as stand-alone consumable services and as building blocks that can be integrated into other services through software interfaces and integration. All of these services need to present in a manner as automated and/or self-service as possible. Seamlessly integrated network services delivered to users, administrators and managers will benefit both community service providers and consumers through tailored value-added service offerings, and enhanced usability and manageability, respectively.

- **Optical Spectrum and Wavelength Services (Advanced Layer 1 Services):** Open Optical Spectrum and Wavelengths upon which regional, national and global infrastructures can be built. Capabilities include colorless capacity for wavelengths greater than 100G and possibly directionless capabilities for rerouting on demand. The ability to carry wavelengths generated by community collaborators and the API's to extend open optical collaborations will be supported.[4]
- **Packet Platform (Advanced Layer 2 & 3 Services):** Packet Platform services that deliver efficient on-demand layer 2 packet networks that support research, flexible academic and production administrative requirements. This service will also offer the traditional core Internet2 IP internetworking that delivers both internet services like Global R&E

[3] One of the planning areas that the regionals and Internet2 must tackle is how to federate automation and software tools in a way that continues to enable local innovation, but also allows for the development of end-to-end services and production quality support tools.

[4] High quality optical services underpin key partnerships like the Energy Sciences Network (ESNet) and CENIC/PNWGP capacity sharing agreement and have additional major stakeholder interests to be integrated in to planning.

Reachability and Commercial Peering, but which will also increasingly include additional private IP networks like shared Cloud Exchange, National Research Platform, LHC-ONE, and other permanent or on-demand virtual networks. An explicit new feature within this service would be private campus cloud interconnection and support for other customized virtual networks across the regional and national networks.

- **Programmable Research Overlay:** Internet2 and regionals should support NSF-Funded (and other agencies) overlay networks that offer a programmable and reconfigurable testbed as a "side path" to the production network to support network research.
- **Custom Platforms**: Mission-built virtual or physical networks built for a specific use case and operated by Internet2 and partners to support custom applications.
- **Build-Your-Own "BYO" Colocation Services**: Support for overlay networks to be hosted in Internet2 operated facilities: GENI, NG Internet protocol, etc.

Common Cross-Cutting Features of all Services

Each of these five core services will include a number of inherent features customized to the technology layer and service, but available through a common framework of self-service automation that allows end users, collaborators, campus IT professionals and regional networks to provision, measure and manage the services as if they had direct control of the virtualized environment. All of this should include a focus on end-user accessibility and usability[5]. This new expectation will require end-to-end implementation of services and new levels of coordination between national and regional networking service providers. While all of the features and integration may not be possible at once, there are some important aspects that can be prioritized, and the rest can be staged for a later deployment.

Research Supporting – Each of the services must envision the researcher use case and support a simple, end to end workflow that allows researchers to accomplish their tasks in an efficient and results oriented manner. This means the technology solutions chosen by Internet2 are only part of the solution. Establishing and communicating best practices throughout the ecosystem and maintaining facilitated support for research solutions must be as core as the simplicity and capability of the underlying technologies. The facilitated support function includes integration with support, communications and materials tailored for researchers beyond the technology platform.

> "The ideal goal is a consistent, scalable, and simple national connectivity approach for provisioning and supporting the complex and evolving specialized needs of collaborative research across the R&E institutions and shared Cyberinfrastructure facilities." – Rick Tutthill and Chris Misra, UMass Amherst, RECINNS paper

[5] It is important to focus the ease of use aspect on end users such as campus CI facilitators or research teams and to use the planning process to engage those individuals directly about what features meet their needs.

Automation – Automation can reduce the human support time required by the infrastructure and increase the usability for non-network and IT communities. Almost all of the features available in each service should be available through a self-management portal that subscribers to the services can use themselves or offer to their downstream users[6]. Configuration of most services and features, telemetry, security and management functions should all be automated to the extent that service requests that do not require physical deployments of new infrastructure should mostly not require human intervention by Internet2 for final configuration. Instead, regional networks and other collaborations should be empowered to leverage their stake in Internet2 directly and Internet2 should be available for consulting and support of advanced features. Features, telemetry and management should explicitly be available through API's in addition to GUI interfaces.

> "replace the mundane, mechanical aspects of network operations with an engine that learns how best to run the network to maximize given objective functions. Here, "running the network" includes configuration, troubleshooting, root cause analysis, corrective actions, fine-tuning and optimization, monitoring, report generation, and other activities typically undertaken by a networks operations crew."- Kareeti Kompella – Juniper – RECINNS Paper

> "Having access to a programmable WAN is important not only for testbeds but also for a number of emerging applications." – Joe Mambretti, MREN – RECINNS Paper

> "Developing a programmability model for the core network to also become more flexible, allowing relinquishing of control over the network substrate to the end users and going beyond simple bandwidth on demand API." – Mark Johnson & Ilya Baldin – RECINNS Paper

Cloud & Peer Connect – Each of the core services will be configured for use as an access channel to commercial & community cloud and other peer networks. Whether by optical wavelength, layer2 packet service or layer3 private IP, a custom network design or a BYO research project, each of the services will have full reachability to the major cloud providers and commercial peering points.

> "support for access to the cloud is necessary. These services are not just for academic enterprises, but also for researchers." – Michael Lambert, PSCC, RECINNS Paper

[6] Each service provider in the end-to-end service chain may have a self-service portal of their own – the manner in which they coordinate (via federation, meta-portal, etc.) to obtain the goal of ease of use for subscribers is TBD.

Embedded Security Services— (physical, embedded, virtualize) – Internet2 will offer a set of security services that secure the infrastructure services and are available on their own as consumable services that can be added to each regional, campus or virtual network. The services need to integrate the community's investment in Trust and Identity tools to the maximum extent possible. Basic security features available as part of Internet2's services should include management of the physical security of the network devices and monitoring and proactive identification of security risks on Optical, Packet and Private IP services. Additional optional security features will be available to support virtual network instances or other emerging requirements.

> "we need infrastructure that includes trust and identity as native components and we need security, including resiliency to ensure that work is not disrupted, pirated or polluted." – Chris Sedore, Nysernet – RECINNS Paper

Regional Integration – Each of these core services will be defined in such a way that a regional network has options to extend the service to an Internet2 member at low cost (via VLAN or wave) without major upgrades to their own services or will be able to replicate and extend the service as part of a more holistic integration of the service offering on to their own. In either model, the technology, the supporting software, and the customer facing capabilities should be seamlessly available to regionals and their end users. As the networking and user communities have greatly benefited from the adoption of technological standards, we again should be looking to adopt emerging public standards describing the delivery of network services over virtualized infrastructures and across network domains. This will help ensure service interoperability within the community and with commercial providers, enhancing the value and utility of community-based services offered to our user community.

Infrastructure Sharing – The core services should be defined to take advantage of potential infrastructure sharing across regional and global collaborators. Through coordination of common service level objectives and long-term agreements, we should allow for consolidation of optical, switch/router, measurement and other infrastructure components that lead to operational savings and reinvestment in new services.

> "There can be no sacred cows. We must challenge the assumption that network ownership [must remain] as we currently understand it and [allow for it to] give way to new models." – Paul Schopis – OARnet - RECINNS Paper

> "These are challenging economic times for every institution and finding ways that we can collaborate to leverage strengths across the community in pursuit of our common goals will help us provide great services, control costs, and keep us all focused on the important elements of our respective portfolios." – Chris Sedore – Nysernet – RECINNS Paper

Global Reachability – Each of the services should be extensible globally at some level, to allow global collaborators to join their domestic colleagues in collaborations. For each service, a framework for extension, management and support of the service should be defined and supported. (The same VLAN and wave extension model that allows low cost extension of Internet2 services through a regional to a campus that is important to regionals should also apply here, allowing global extension of US-based services without requiring upgrades to current global infrastructure.) Internet2 should work with the NSF and the IRNC community to grow support for International transit across the Internet2 backbone.

Measurement Services - (performance, operational telemetry, logging) – To support integration into other environments, all of the operational telemetry normally available to the hardware owner/operator must be virtualized and available for each service to its users as if it was their own hardware. This means telemetry ranging from port interface utilization and errors, to system logging, to flow data should be available in common formats to all service subscribers. Similarly, performance monitoring tools must be virtualized and available to measure each common and/or virtual network at each layer – as well as being accessible to VOs that want to gather and correlate with other levels of performance data (system, application, etc.).

Research Data Services - All of the measurement services data should be made available, through appropriate privacy protecting data-sharing agreements to researchers with valid interests in accessing data from live operating environments.

Implementation Expectations and Service Economics:

As these services and their features are defined, there are also expectations around economics, operational coordination and other aspects of the services that must be managed to meet business and other expectations. To be clear, in all of the services, it is expected that the community will find efficiencies that both improve our services to our core constituencies and allow us to offer more capacity and capability. ***Pressure to reduce fees to regionals as well as to increase service levels and capacity must be addressed.*** Further to the goals of improving cost efficiencies, service delivery and manageability should also be enhanced through service integration and orchestration. Improving the user experience with network services is gaining in importance as more of our community are engaging directly with commercial providers who have an enterprise-level user experience model that will drive expectations.

> From an ecosystem perspective, it is important to optimize our collective investment and not spend more on services or architectures than is necessary and prudent. Our national interests are best served when Internet2 AND the regional networks are successful. - From Guiding Principles created at 2016 Tempe Regional Principal's Meeting

While these expectations are cross-cutting and affect Internet2 and its regional partners equally, we can be explicit about expectations for Internet2's infrastructure program to include:

- Investment in more efficient technologies and software automation, will provide substantial *recurring operational savings* across the current total annual expense for maintenance, power, colocation, and operations staffing. **Such savings are necessary to enable the new cycle of investment and the ongoing capacity growth that the network will require as well as address pressure to reduce the cost per service delivered.** The planning and procurement teams will be explicitly tasked with finding these savings through negotiation and design considerations.

  Example: Internet2 currently spends about $15.4M per year on maintenance for its Optical, Route/Switch, fiber maintenance, Operations Support contracts, power and colocation. As part of the network investment, staff will be asked to find savings of 20 % or more across these expense categories to support the expected expanded capacity and services costs and potential changes to the fee structure.

- Savings realized through reinvestment must allow Internet2 to *offer a substantial increase in the bandwidth* by 2021/22 offered to Internet2 connectors, enabling full 100G peering, dedicated research 100G and other features. In addition to funding necessary capacity growth, Internet2, Regionals and Campuses will need to discuss potential changes to the service fees across the ecosystem.

  Example: When cost savings can be realized from technology re-implementation, automation and other activities, Internet2 should work with the community to reconsider how services are offered and how campuses and regionals wish to consume that new capacity. This may result in the need for a community business model review across all of the service stakeholders.

Internet2 Advanced Layer 1 Services

Service Definition: 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 400 Gigabit. There are more than 50 locations on the 16,000-mile fiber footprint where waves can originate or terminate. Participants can procure a single wave from one point to another, or an entire nationwide network. 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, 100 and 400 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 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 2019. There are more than 50 locations on the 16,000-mile fiber footprint where waves can

originate or terminate. Participants can use the Spectrum Service for a single wave from one point to another, for a group of waves or for an entire nationwide network.

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 as used by Internet2 to light the Optical System or equipment from other manufacturers with prior approval.

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. 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 (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 in to 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 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.

Service Definition: Dark Fiber Service

Internet2 Dark Fiber Service provides custom, cost-effective solutions for any dark fiber requirement – acquiring, holding, and distributing dark fiber network assets—subject to availability by dark fiber vendors.

Requirements for Participation in this Service

To participate in this service a member organization will need to be in a position to light the fiber asset when it is delivered.

Service Characteristics

Dark Fiber Assets typically can be acquired for both metropolitan and long-haul applications. Internet2 can also assist in IRU renewals for next generation Ultra and ULL fibers on older IRU's. As part of the service, Internet2 typically identifies available providers for the member's routes of interest, consults with the member about available paths and negotiates R&E community-friendly IRU and Operations and Maintenance agreements with the providers. Internet2 can hold that fiber asset and relicense it on a long-term basis to a member or facilitate a contract between the underlying provider and the member organization.

Internet2 AL2S Service

Service Definition: Advanced Layer 2 Service

The Internet2 Advanced Layer 2 Service (AL2S) allows users connected to Internet2 at 10G or 100G to create their own VLANs on the Internet2 AL2S backbone. Static or Dynamic, point-to-point or multipoint, intra-domain or inter-domain, AL2S puts control of the backbone VLANs into the users' hands for the creation of purpose-built private networks using Internet2 infrastructure.

Using the Cloud Connect Portal/OESS provisioning tool (https://www.internet2.edu/products-services/advanced-networking/oess/) users can provision networks, establish automatic network failover, define per-interface access permissions, and generate automatic per-VLAN statistics. It includes a simple and user-friendly web-based user interface as well as several programmable API's.

Requirements for Participation in this service

This service is available to Internet2 members and typically is used by Internet2 Connectors. When an AL2S connection is established, the users of that connection will be able to create accounts in CCP/OESS as well as create user groups indicating who else within AL2S Is allowed to create connections to that organization. There are no limitations to the number of VLANs that may be created on a given connection.

Service Characteristics

AL2S is a highly available, On-Demand and User-Configured service. Current access types include 10G and 100G ethernet. The service is provisioned as label switched encapsulated paths across the Internet2 network. Provisioning is accomplished through the CCP/OESS portal and users and groups may configure VLANs to any other AL2S participant that is part of their group.

Connections over the Internet2 network using AL2S will have automatic failover based on label switch paths. Should a fiber cut or other failure cause a path to be unavailable traffic will move to the secondary path and notification provided. It is also possible to restrict the connection to a primary path only in the case where there are specific characteristics of a path required.

In addition to the web portal, AL2S also supports several API's for direct programming. These include REST-API, GENI Aggregate manager and NSI.

New features that could be introduced as part of NGI include (pending demand and interest):

- Custom telemetry streamed to each network owner for each network they have provisioned on the infrastructure.
- White-Labeling of the Cloud Connect Portal / OESS with regional and/or campus member branding.
- 400G Ethernet access
- Interoperation of CCP/OESS with regional networks to allow end-to-end provisioning.

Internet2 AL3S Services

Internet2 Advanced Layer 3 Services include the traditional Internet2 R&E Service, the Cloud Exchange Service (currently known as TR-CPS) and customer layer-3 services such as Cloud Connect, LHCONE, etc. For these service definitions, we are describing each of these sub-services, available on an AL3S connection, separately. As we consider future business models, it is possible these individual services could be developed differently and be offered with differentiated service fees separate from a standard AL3S connection fee.

Service Definition: AL3S Connections

The Internet2 Advanced Layer 3 Service (AL3S) service allows users connected to Internet2 at 10G or 100G to participate in one of several national and global Layer 3 networks. AL3S is currently provisioned using multiple 100G links and includes headroom policies, security, measurement, analytics and performance characteristics that make it unique for R&E. Current offerings include 10G and 100G access to the AL3S network at over 50 access points across the US.

Requirements for Participation in this Service

This service is available to Internet2 members and typically is used by Internet2 Connectors.

Service Characteristics

AL3S is a highly available nationwide layer 3 network. Each Layer 3 network provided on AL3S is provisioned as label switched virtual network that can be carried across the Internet2 network. IP networks on AL3S have automatic failover based on label switch paths. Should a fiber cut or other failure cause a path to be unavailable, traffic will move to the secondary path and notification provided. Provisioning is accomplished both through the Cloud Connect/OESS portal and through ticket requests to the NOC. perfSONAR nodes are available throughout the network for testing of performance. The network is interconnected Internationally to similar R&E networks in over 100 countries.

New features that could be introduced as part of NGI include (pending demand and interest):

- Custom telemetry streamed to each network owner for each network they have provisioned on the infrastructure.
- Additional provisioning and control of AL3S networks in the Cloud Connect Portal / OESS with regional and/or campus member branding.
- Additional security and “clean pipe” services.
- 400G Ethernet access.
- Interoperation of CCP/OESS with regional networks to allow end-to-end provisioning.

Service Definition: R&E IP Service

The Internet2 R&E IP service is available to all members of Internet2 who follow Participation policies for Internet2. This service provides Participants access to a subset of the global Internet’s IP routes, and thus connectivity through the Internet2 network, for the following sets of participating organizations in the Internet2 network:

- All Connector organizations (regional and state networks) including the educational, library, medical and other institutions connected to those state and regional networks.
- All connected Internet2 member Higher Education Institutions.
- All International R&E organizations with peering relationships to Internet2.
- Connected research arms of Corporate members of Internet2.
- Federal Networks connected to or peering with Internet2 including ESnet and NOAA, for example.
- Connected Research Instruments (LHC, LSST and SKA for example).

Requirements for Participation in this Service

Both Internet2 membership and an agreement that allows participation is required before passing traffic on the R&E IP Service. Following appropriate participation agreements, an organization must have a path to an Internet2 router.

It is currently requested that participants in this service also adhere to the rules set out in MANRS (Mutually Agreed Norms for Routing Security). More information on MANRS can be found at MANRS.org. While not a requirement for participation in the R&E service it is strongly encouraged.

Service Characteristics

The R&E IP Service is a base-level service running on the Internet2 backbone network. As such it is a highly available, resilient and reliable service. A connector's experience with this service will be dependent on the nature of their connection to the Internet2 backbone. Those connectors who have dual connectivity to the backbone should expect performance in the range of 99.999% for the service as a whole.

After an organization is connected a BGP session between the two routers will be set up and the connecting organization will begin to receive the R&E service. They will also be able to inject routes into the service ensuring that other participants in the service will also have access to them.

Service Definition: Cloud Exchange Service

The Internet2 Cloud Exchange IP service is available to connectors and participants of Internet2. This service provides participants access to a variety of content providers and other commercial organizations at the following peering locations across the U.S.: Seattle, Los Angeles, Chicago, Dallas, Ashburn and New York. Cloud Exchange provides high performance, low latency, and efficient (1 hop) access to some of the top content destinations in the world including: Google, Akamai, Netflix, Facebook and other commercial content providers. Cloud Exchange also provides access to the public portion of key cloud providers such as Amazon, Google and Microsoft. Participants in Internet2 NET+ program are also reachable via the Cloud Exchange. The service supports IPv4, IPv6 and multicast.

In the context of this service “peering” is defined as a settlement free exchange of routes with another organization for the purpose of exchanging traffic.

Requirements for Participation in this Service

To participate in this service an organization must have a connection on an Internet2 router. After an organization is connected, a BGP session between the connector router and the Internet2 Cloud Exchange VRF (Virtual Routing and Forwarding) instance will be set up and the connecting organization will begin to receive the Cloud Exchange service. They will also be able to inject routes into the service. These routes will be passed on to the Cloud Exchange peers though Internet2.

It is currently requested that participants in this service also adhere to the rules set out in MANRS (Mutually Agreed Norms for Routing Security). More information on MANRS can be found at MANRS.org. While not a requirement for participation in the R&E service it is strongly encouraged.

Service Characteristics

The Cloud Exchange Service is a base-level service running on the Internet2 Backbone network. As such it is a highly available, resilient and reliable service. A connector’s experience with this service will be dependent on the nature of the connection to the Internet2 backbone. Many peers connect to Internet2 at multiple locations. This provides added reliability for those organizations.

Participants in this service can give local preference to the Internet2 Cloud Exchange routes to ensure that the Cloud Exchange paths will be prioritized over other Internet connections for outbound traffic. Internet2 strives to provide the best path to the peering locations but it can happen that a provider sees a non-Internet2 path as preferable for return traffic. Thus, Internet2 cannot guarantee that the Internet2 path will always be the return path selected.

At present this service is manually established. Organizations wishing to participate in the service are able to request it through the Internet2 NOC. A longer-term goal for this service is to have service provisioning be automated.

Service Definition: Cloud Connect Service

The Internet2 Cloud Connect Service provides direct private peering to three major Infrastructure as a Service (IaaS) cloud service providers leveraging the regional and national network infrastructure. The service leverages both the AL2S and AL3S service offerings to provide connectivity to Amazon Web Services’ Direct Connect, Google Cloud Platform’s Dedicated Interconnect (DX) and Microsoft Azure ExpressRoute at speeds up to 10G. Both

Layer 2 and Layer 3 service is offered in three cities: Ashburn, VA, Chicago, IL, and Dallas, TX with Sunnyvale, CA to be added soon. Currently each location provides 10G connectivity to AWS, with Google and Microsoft employing redundant 10G connections.

Requirements for Participation in this Service

Internet2 members may use this connectivity service through an Internet2 Connector. Organizations must establish their own subscriptions to the private "direct connect" services with each of the cloud providers to obtain access to IaaS resources within the provider's cloud environment. Subscription rates to the cloud providers vary depending on the cloud provider chosen. Regional networks are encouraged to pre-provision vlans and establish user groups in the portal to likely user/members of this service to the Internet2 network.

Service Characteristics

This service is available as an on-demand user configured service using the Cloud Connect Portal CCP/OESS. Users can provision Layer 2 or Layer 3 network paths, establish automatic network failover, provide per-interface access permissions, and generate automatic per-VLAN statistics. CCP/OESS includes a simple and user-friendly web-based user interface as well as a web services API. The CCP/OESS web tool also allows configuration of Layer 3 characteristics and integrates with Amazon and GCP API's.

Once established these private connections will run across the Internet2 backbone network, a highly reliable and resilient network. The connector can implement a design that ensures resiliency between the Internet2 backbone and its members' networks, such as ensuring the VRF is enabled on each of its Internet2 connections.

New features that could be introduced as part of NGI include (pending demand and interest):

- Custom telemetry streamed to each network owner for each network they have provisioned on the infrastructure.
- White-Labeling of the Cloud Connect Portal / OESS with regional and/or campus member branding.
- 100G cloud connectivity
- Interoperation of CCP/OESS with regional networks to allow end-to-end provisioning.
- "Clean pipe" security support for the transit service, possibly including intrusion detection, DDoS avoidance, host scanning/reporting and anomaly detection.
- On demand perfSONAR nodes at the cloud-interconnection points

Service Definition: Full Transit IP Service

In response to a small number of Internet2 members and Connectors that wish to leverage their Internet2 connection investment to access commodity resources, Internet2 will continue

to offer Internet2 Full Transit IP service. Internet2 will provide access to one or more Internet Service Providers (ISPs) enabling access to a full Internet routing table via a full transit VRF. Participants will specify the amount of bandwidth they will require.

Requirements for Participation in this Service

To participate in this service an organization must be connected to an Internet2 router, either directly or through a state or regional network, the latter using a VLAN established by the state or regional network to the Internet2 connection point. An organization wishing to participate will create a BGP session with the Full Transit VRF and receive the full routing table. The device used for this BGP session will need to be of sufficient capacity to carry all of the routes, both v4 and v6, that they will receive from this BGP session.

Service Characteristics

The Full Transit service will be a full production level service within Internet2. It will be architected to eliminate single points of failure and to maximize resilience throughout the network. The amount of bandwidth maintained by Internet2 will be commensurate with the service demand from participants. Access to the Internet2 analytics service is included.

Participants will specify their bandwidth requirements. Burst capability beyond that amount will be allowed.

New features that could be introduce as part of NGI include:

- “Clean pipe” security support for the transit service, possibly including intrusion detection, DDoS avoidance, host scanning/reporting and anomaly detection.

Service Definition: Internet2 L3 VPN IP Service

This service provides participants the ability to establish a private Internet Protocol (IP) network across the Internet2 infrastructure. A separate Virtual Routing and Forwarding instance (VRF) is established for each private network. Access to this VRF is by invitation of the group requesting the private network. There are no restrictions on the number of organizations that participate in such a private network. Example users of this service would include LHC (LHCONE), DDoS Volumetric Mitigation Service and the Open Science Grid. A private network may also be established by a group of universities or departments with a shared interest in a research or other topic. Connectivity to various cloud providers such as AWS or Google is also be available.

Requirements for Participation in this Service

The Internet2 L3 VPN IP service is available to connectors and members of Internet2. To participate in this service an organization must be connected to the Internet2 network and

have accounts on the Internet2 Cloud Connect portal. The VRF for the private network can be extended to any number of devices on the Internet2 infrastructure. Regionals are encouraged to pre-provision extensions from the Internet2 network edge to potential users in advance.

When an organization wishes to join one of these private networks a BGP session between the member and the appropriate VRF will be established (typically over a regional VLAN or regional's own VRF). Then the organization will be connected and will be able to inject its routes into the VRF insuring that other participants in the private network will be able to connect to them.

Some private networks may choose to require participation or partial participation in MANRS (Mutually Agreed Norms for Routing Security). More information on MANRS can be found at MANRS.org

Service Characteristics

Each L3 VPN will be routed on the Internet2 AL3S backbone network, benefiting from the reliable and resilient backbone and its operational policies. Regionals can enhance this capability by implementing a design that ensures resiliency between the Internet2 backbone and its members' networks, such as ensuring the VRF is enabled on each of its Internet2 connections.

This service will be available as an on-demand, user configured service using the Cloud Connect Portal/OESS. With CCP/OESS (https://www.internet2.edu/products-services/advanced-networking/oess/) users can provision VPNs, establish automatic VPN failover, provide per-interface access permissions, and generate automatic per-VLAN statistics. It includes a simple and user-friendly web-based user interface as well as a web services API.

New features that could be introduce as part of NGI include:

- Custom telemetry streamed to the VRF owner for each VRF and its infrastructure.
- "Clean pipe" security support within the VRF, possibly including intrusion detection, DDoS avoidance, host scanning/reporting and anomaly detection.
- White-Labeling of the Cloud Connect Portal / OESS with regional and/or campus member branding.
- Interoperation of CCP/OESS with regional networks to allow end-to-end provisioning.

Potential NGI New Services

Service Definition: National Science Network Service

The Internet2 Science Network Service would be available as a service within AL3S to members of Internet2 and key science collaborators. This service would provide participants the ability to join a dedicated private Internet Protocol (IP) network across the Internet2 infrastructure for science instruments and Data Transfer Nodes (DTN's). The Science Network Service will be an environment that is a purpose-built, **production** service, not a testbed. It would provide a high performing production environment in support of Science Workflows. By tailoring and developing a complete environment for science, the new environment can be more tightly coupled to support the emerging distributed scientific community. While the Science Network Environment would not be a physically separate network from the Internet2 R&E Network, it could have different characteristics related to security, capacity, telemetry, etc.

Requirements for Participation in this Service

There are a variety of ways research teams could participate in this service. However, participation is limited to portions of the institution engaged in sending/receiving scientific workflows. This may be through use of a Science DMZ or something similar. The institution's entire IP network prefix lists would not be allowed to be advertised to the Science Network.

A separate virtual or physical connection to Internet2, typically via a regional network, will be necessary but there will be multiple options for creating this connection. Connections can be built using optical devices, switching technologies or routing devices.

Service Characteristics

This Service would be tailored to make the end-to-end experience for the scientist 1.) easy to use, 2.) more transparent, 3.) high quality, 4.) productive. Key aspects likely include science-focused telemetry, enhanced security protections, analytics and reporting appropriate for grant programs, researcher directory integration and other research-centric capabilities.

The infrastructure established for the purpose of supporting science will support dynamic changes as necessary to meet the need of the scientist. The environment will:

1) Provide access to the many communities and tools already in use within the scientific community, including work by Internet2 to link network programmability to commonly used science applications.

2) Allocation of resources to adapt to meet the varying and changeable needs of the communities using it.
3) Be able to incorporate new tools and communities of interest as they develop.

The service will be instrumented sufficiently to allow information to be provided to funding agencies and other interested parties with a stake in the usage, usability and performance of the service.

To protect the integrity of the data, security features to safeguard the storage and transmission of research information will be integrated into the production environment.

Rapid Deploy Layer 2 Private Network Service

Service Definition: Rapid Deploy Layer 2 Private Network Service

The Rapid Deploy Layer 2 Private Network Service is intended to allow connectors to extend to remote peering points without needing to invest in their own connectivity, colocation or equipment. A key aspect of the service is to pre-provision a pool of capacity so that a requesting entity could start working on actually provisioning the new connection within hours of requesting it. The service would allow Internet2 connectors to obtain one or more 10G connections delivered on the same day as requested via a "one click" service schedule amendment. The intention is to make this extremely easy by pre-provisioning the port assignment, LOA for a cross connect, and authorization to configure the connection using existing agreements.

Requirements for Participation in this Service

To participate in this service an organization must be an Internet2 connector with a current AL2S or AL3S agreement.

Service Characteristics

This service is envisioned to be available as an on-demand user configured service, likely in CCP/OESS. It is intended to be complementary to the new Virtual Router and Shared Router offerings as well as AL2S, AL3S and Cloud Connect options. Within CCP/OESS users will be able to provision networks between locations on the network, establish automatic network resiliency, provide per-interface access permissions, and generate automatic per-VLAN statistics.

VLANs and trunks established from connectors to these service connections would be protected on the Internet2 backbone over resilient paths by default, normally taking the shortest path and automatically restoring on to other available paths during an outage (such as a fiber cut.) Headroom for these services would be managed as part of the aggregate traffic headroom

management for Internet2's IP backbone. Optionally, this service will be available in a premium configuration that will leverage AL1S waves or QoS-protected AL2S network capacity.

To enable this service, Internet2 will:

- pre-allocate a pool of 10G connection on its network devices in key peering locations,
- extend those connection to a patch panel, and
- prepare (in advance) letters of authority (LoA) for interconnection.

Information regarding the allocated 10G connections will also be loaded in the Cloud Connect Portal/OESS, allowing connectors to use the connections as soon the agreement and a cross connect to the connection are completed.

New features that we could also introduce as part of NGI (pending input and demand)

- Custom telemetry streamed to the member for each provisioned connection.
- Virtual perfSONAR nodes at the peering point.
- Premium bandwidth protection back to the connector.
- "Clean pipe" security support within the VRF, possibly including intrusion detection, DDoS avoidance, host scanning/reporting and anomaly detection.
- White-Labeling of the Cloud Connect Portal / OESS with regional and/or campus member branding.
- Interoperation of CCP/OESS with regional networks to allow end-to-end provisioning.

Internet2 Virtual Cloud Router Service

Service Definition: Internet2 Virtual Cloud Router Service

Internet2 is considering a new service, the Virtual Cloud Router Service, in support of community requests for cloud and commercial peer connectivity to establish private infrastructure into metro peering regions. This service provides participants the ability to create and configure a virtual router as part of their own network at a remote peering point. This service may be used in conjunction with other Internet2 services, such as AL2/3S, AL1S Waves, Rapid Deploy Layer2 PNI, to establish layer3 connectivity in metro peering regions and between cloud service providers.

This service would facilitate connectivity between resources as part of campus' own network without needing to contract with service providers for colocation (rack and power), connectivity, and without needing to procure and deploy hardware. The process of an organization procuring hardware, colocation, connectivity then deploying assets is resource extensive and often protracted and complex. Internet2's Network and Virtual Cloud Router and

Cloud Connect services would eliminate these challenges with a software-based slice of a server in the metro area.

Requirements for Participation in this service

Internet2 members can acquire this service in a peering point. There is a requirement to use Internet2's AL2S and AL3S network to create network capacity between the cloud router and the member's home campus as well as the various service providers in one or more peering points.

For regional network providers, it is anticipated that Internet2 and the regional would establish a white-label capability to allow the regional to use Internet2 supported portal, management and provisioning tools within their own organization's web presence.

Service Characteristics

The service would be initially available in major peering points. Based on demand, additional nodes could be added. The service allows for on-demand rapid provisioning, multi-cloud inter-connectivity, multi-region access and pay only for what you need.

- Infrastructure without capital investment; member can "log in" to its virtual remote router and configure it as if it is their own.
- Ability for member to use its own BGP and internal routing protocols
- Interconnected with Internet2 Layer 2 and 3 services (Cloud Exchange, R&E IP)
- Integrated with Rapid PNI service to allow physical connections to be provisioned in hours when necessary.
- On-Demand: Rapid turn-up / turn-down
- Secure: Any-to-Any Secure Environment between clouds;
- Localized Routing: No hair pinning traffic between clouds (higher performance)

Potential Virtual Routing Zones:

US East	US Central	US West
New York Ashburn	Chicago Cermak Dallas InfoMart	Los Angeles San Jose SV5 Seattle

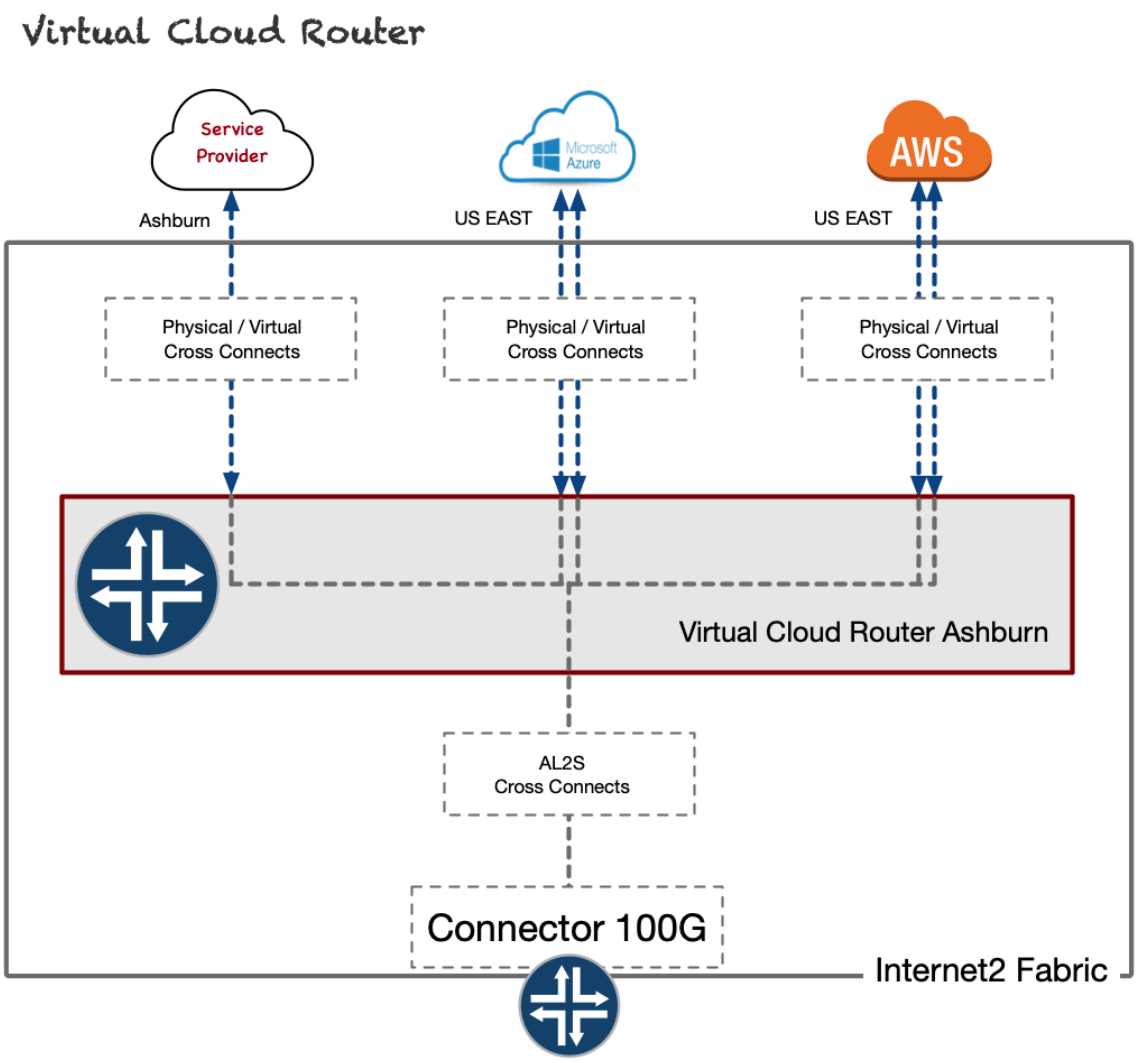

Internet2 Remote Dedicated Router Service

Service Definition: Internet2 Remote Shared Router Service

Internet2 would introduce a new service in support of requests to share infrastructure. This service is envisioned to facilitate member's private infrastructure and connectivity needs in areas outside their operating area without needing to contract with service providers for colocation (rack and power), connectivity, and without needing to procure and deploy hardware. The service could be based on "slicing" a larger Internet2 switch/router or could include allocation of a physical small form-factor physical-router to a member. In either case, the member would have complete configuration control of the dedicated hardware. Access to

Internet2 AL1S, AL2S, AL3S and cloud resources would be enabled. Workflow would be designed for rapid deployment and reconfiguration.

Requirements for Participation in this service

This is envisioned to be an independent service available to Connectors and International partners. There would be no requirement for interconnection with Internet2's network or other services to take advantage of this service, however Internet2 would facilitate an AL2S/3S connection at 10G to each dedicated router at no additional cost if desired.

Service Characteristics

This dedicated router provides hardware isolation of a connector or international partner's network configuration while allowing sharing of facilities, staff and support costs for private infrastructure in key facilities. Dedicated hardware would be provided to each service user and the service user would be expected to configure and manage its dedicated hardware.

- Infrastructure without capital investment; member can "log in" to its physical remote router and configure it as if it is their own
- Ability for member to use its own BGP and internal routing protocols
- Interconnected with Internet2 Layer 1, 2 and 3 services (waves, AL2S paths, Cloud Exchange, R&E IP)
- Member can add additional physical connections on demand within the capabilities of the card slot or physical device without additional fees.
- On-Demand: Rapid turn-up / turn-down of additional peers or connections; workflow pre-optimized for "same day" progress
- Localized Routing: No hair pinning traffic between clouds (higher performance)
- Internet2 will monitor the hardware, provide break-fix of any failed hardware components and facilitate remote hands support. Support for break-fix would be on a 4-hour hardware replacement
- Internet2 can provide maintenance notifications and support the base chassis
- Internet2 will provide LOAs for ordering cross connects at the time any hardware is added to the router, allowing immediate provisioning of additional services.

There are two versions of this service that could be delivered, pending interest and planned RFP and community evaluation during 2019 of new Internet2 platforms for NGI:

1. Dedicated Small Form Factor Router – This might look something like a dedicated Juniper 204 router. This is a 1 rack-unit sized router with 8 10G and 4 100G connections at a low cost. One or more of these devices could be dedicated to an Internet2 member and made accessible via the Internet2 management network. A 10G interconnection to the Internet2 network could also be arranged, with 100G provided for a premium fee.

2. Juniper Node Slicing (JNS) technology - Juniper Node Slicing (JNS) can be thought of as the next generation of virtualization. JNS completely separates the routing engine and line card resources between different partitions, called Guest Network Functions (GNFs), within a single large chassis. These GNF's can then be allocated to tenants as required. The GNFs and associated line cards share the same switch fabric. Isolation is possible by instructing the fabric to only interconnect certain components.

Potential Remote Dedicated Router Sites:

US East	US Central	US West
New York 111 8th or 32 AofA Ashburn	Chicago Cermak Dallas InfoMart	Los Angeles 1 Wilshire San Jose SV5 Seattle

DRAFT for discussion with Principals

POTENTIAL NEXT GENERATION INFRASTRUCTURE SERVICE MODELS

The Research and Education networking community is being asked to address several issues facing all of us with the implementation of the Next Generation Infrastructure. These include:

- Expanding capacity offered to regionals, including removal of Cloud Exchange (formerly TR-CPS) utilization caps, introduction of private cloud access, and expansion of the capacity for research;
- Emerging research-support opportunities like implementation of infrastructure and programs for a national research platform;
- Introduction of automation, self-service tools, and API's, all to enable the sharing of community infrastructure and ease-of-use for users of end to end services; and
- Delivery of these features and improvements with the goal to lower total cost to operate the network for the services that it delivers.

We feel that now is a good time to consider new Service Models for delivering current services and new services that support the key themes of flexibility, programmability and agility. The Service Models endeavor to move us away from our seven-year-old port-based model, to a platform that has the ability to meet the requirements our members told us they need to be successful.

Each of the models is built on the assumption that the connector community will agree to make investments similar to those they make today for Internet2 services, and that Internet2 will provide more value and new functionality responsive to recent requirements articulated by the community for that investment.

Model 1: Provide more capacity for the current services at current fees (More for your Money)
Continue to offer current model & offering additional service capacity. For example, offer 2x 100G service capacity for the cost of 1x100G today and 4x 100G service capacity for the cost of 2x100G today. We could include in this offering 100G services for Cloud Exchange (formerly TR-CPS) and 10G service to Cloud Connect.

Model 2: Base Connector Fee with smaller Services Fees
This model recognizes the common contributions necessary to sustain a high quality, reliable set of national services not so much through service based fees, but through a standardized per-connector fee. This model covers most of the base costs for Internet2 network operations through these base fees and then recovers the remainder through service fees. This approach would make operations more consistent and less subject to questions about who does and who does not participate in various services.

One way to implement this model might be to determine a higher Network Membership fee that would provide some level of base service. Then a smaller fee would be charged for the individual incremental services.

Model 3: Service Menu model

In this model, Connectors would receive a credit equal to the amount they currently pay in fees for standard services to be used to procure any connectivity service that is available through the Internet2 Network Services (NS) Service Catalog at the listed fees.

Connectors could purchase anything in the Service Catalog with a value up to their current subscription fee (e.g., Connector A procures= $400K in AL2/3S Services today, so they could obtain up to $400K of any of the re-priced services from across the entire catalog). Additional services, beyond the current investment, would also be available.

Additional Considerations

- **Service Menu Pricing**: we may also want to consider the way we determine the fees for the items in the services being offered. We should consider that the fees that a Connector would pay for services may be lower than fees that we would charge other members, such as government agencies or industry members. Such a model would recognize the investment commitment that the Connectors have made and provide room for Connectors and Network Members to offer these services to their members and recover any costs they may incur.

- **Community Anchor Program (CAP) Fees:** we would also like to have a conversation with connectors and network members about whether they would like to include the revenue paid in current CAP fees as part of the revenue neutral calculation. We would still need to collect for CAP fees that are needed for the programmatic portion of the CAP program separately.