Network Architecture, Operations and Policy Program Advisory Group
Future Peering Service Framework Subcommittee Report
FINAL
March 12, 2014

Contents

Executive Summary

The committee agreed on the following principles as being vital to ensuring that the next evolution of a peering service continues to meet the needs of our community. These principles include:

- Virtualization of the network rather than using a separate infrastructure
- Allowing NET+ services as part of the peering service
- Creating a new Headroom Practice
- Continuing to offer dedicated TR-CPS ports
- Providing at least 24 months assurance of TR-CPS availability before planned major changes
- Increasing transparency and governance
- Addressing dedicated staffing
- Expanding collaborations to include international peering

Strategic Value of the Service

Today, all major Internet Service Providers (ISPs) interconnect with one another and agree to exchange traffic at public and private peering locations. Most of these exchanges are "settlement-free", i.e., the parties pay only their own incremental costs to interconnect one with the other. The Research and Education networking community has an established history of seeking settlement-free peering.

The Internet2 Transit Rail-Commodity Peering Service (TR-CPS) has traditionally brought great value to the community. It was originally established on the notion that unused capacity on the Internet2 backbone represented a sunk cost and could be leveraged to provide additional value for members with little additional expenditure. With the deployment of the current 100 Gbps backbone the community believes it is appropriate to reevaluate, reinvent and reinvigorate the peering service and consider all of the alternatives.

Goals and Objectives

At the highest level, the goals and objectives of the next evolution of peering service for the Internet2 community can be summarized in a straightforward manner. It should be a cost effective, ubiquitous service that is available to all Internet2 constituencies. It must be a flexible service that can be used effectively by regional networks and higher education institution members (i.e. local policies can be accommodated). And it must be a service that provides value for commodity use of the network without jeopardizing the research and educational mission of the community (i.e., it will augment the R&E components and not compete with them).

Architecture
The design of the service should be based on the principle of efficiently integrating the next evolution of the peering service with the existing Internet2 network infrastructure. The design should include considerations for the best utilization of the existing Internet2 backbone routers to implement both the R&E Layer 3 service and the next evolution of the peering service. The design should also be flexible enough to support additional services as needed, such as Net+ or a high performance research service, etc. The committee discussed utilizing separate Virtual Routing and Forwarding (VRF) instances for each service, but elected to defer the decision on actual design to the NTAC working in conjunction with Internet2 staff.

Reliability
In order for the peering service to function as required, reliability is a paramount consideration. The community, including the regional networks and members, expects the service to perform at production quality, with an availability approaching 100%. The expectation of reliability is also required for the service to be a member in good standing of the peering community.

In defining the peering architecture, every effort should be made to ensure isolation between services in such a way that the reliability of the peering service is not impacted by other services. The peering traffic should be protected from competing with other classes of traffic from the connector handoff to the peer handoff so that the bandwidth committed to the customer (currently 20 Gbps of peering traffic per 100 Gbps interface) is appropriately prioritized. The interconnections between the Advanced Layer 3 Service (AL3S) and the peering service should be diverse enough to guarantee a high-level of resiliency in the case of link or node outages.

The peering service should be designed with the capability to grow as the needs of the community expand. The expectation should be that the significant increase in bandwidth demand that the community has seen over the past decade will continue unabated.

The architecture of the service should support flexible levels of reporting and visibility into the service's inner workings. Flexibility is required in order to provide information in the appropriate level of detail to the appropriate audience. Due to the sensitive nature of some aspects of peering, it is important to allow for the peering service operators to finely control the level of detail provided to each audience.

The architecture of the peering service should support greater integration with the regional networking community's suite of services. The idea of a service-oriented networking approach has been discussed widely in the community of late, in an attempt to provide a more integrated end-to-end network service that can support Net+ and similar initiatives.
In this spirit, it should be recognized that connectors might require some flexibility in the peering service definition in order to optimize the blended service they deliver to their customers. The service should accommodate configuration optimizations based on a particular connectors private peering and caching strategies, as an example.

Besides growing within the current geographic bounds, the next evolution of the peering service should be designed to support greater cooperation with international partners, particularly other NRENs that have already established peering services of their own

Operations
The Committee recommended that a Lead Peering Coordinator (LPC) be identified to oversee the service. The LPC role requires that the individual be community-oriented and have the ability to deftly make decisions within the framework defined by a group referred to as the Peering Advisory Subcommittee (PAS). S/he would serve as the face of the peering service to the Internet2 community, as well as to commercial peers with whom s/he would engage. The committee suggests building a staff consisting of a full-time LPC and one or two junior peering coordinators as a transcontinental team.

The committee agreed that operations would be handled by a single entity that would operate under the guidelines provided by a PAS under the auspices of the NAOPpag. The PAS would define policy and peering strategy to provide a framework in which the LPC would work. The committee suggested the creation of a steering group made up of two or three peering practitioners from the community who would serve as a "sounding board" for the LPC. This group would be available to provide community input to the LPC on tactical decisions for which the LPC seeks input.

Financials
Internet2 and CENIC have compiled a shared view of the TR-CPS financials for the past four years. The financial detail provided in this report includes the direct costs of the program, but like other Internet2 network-dependent programs, the generalized use of the Internet2 backbone, staff and overhead are not directly charged to this program.

During 2014, it is expected that substantial capital expansion will occur for TR-CPS, with new higher-density and 100 Gbps capable core switches deployed to at least three TR-CPS peering cities. Consolidation of R&E and TR-CPS offers an opportunity to consider how many routers are necessary in each key R&E and TR-CPS geographic location. While consolidation might provide substantial operating savings over time, it is equally possible that scale or resiliency benefits of having multiple routers in key cities will continue to make sense. We would suggest the committee that will look at architecture consider expense as one of several important factors and that Internet2 be asked to report on potential consolidation cost implications as part of the analysis.

Recommendations
The Future Peering Service Framework Committee recommends that the NAOPpag charge Internet2 staff to engage community participation via the Network Technical Advisory Committee—NTAC in implementing the next generation of peering services for the Internet2 community and our global peers.

20-month goals

- Virtualize the network – the network will leverage the existing 100 Gbps infrastructure and technically stable solutions to implement the service
- NET+ will be offered as a part of the service
- Headroom Practice – A more coherent approach to capacity and costs is to set data driven goals to assess those needs and costs so the community can make informed decisions. A "prime directive" of do no harm is mandatory for all types of traffic
- Flexibility – The service will offer dedicated ports for commodity services as well as the more traditional converged services currently offered
- Predictability – At least 18-24 months notice of any changes to the program's structure will be provided to allow regional networks and member schools to make adjustments to their operations and budgets
- Establish a default Autonomous System Number – The goal is to set a default "Higher Education" AS number to enable content and service providers "easy" access to the HE domain within the United States. (See next section for global ASN)
- Staffing and Oversight --
    - Creation of a Peering Coordination Group consisting of Internet2 dedicated staff--Lead Peering Coordinator and two other staff members
    - Creation of a Peering Advisory Subcommittee (five to seven community members) would define policy and peering strategy to provide a framework in which the Peering Coordination Group would work.
    - Creation of a steering group made up of two or three peering practitioners from the community who would serve as a "sounding board" for the Lead Peering Coordinator.

Three-Five Year goals

- Future network requirements should be more data driven than casual estimates.
- Direct connections to the peering service should continue to be offered,
- Headroom practice should be revisited so that the contradictory requirements of providing service guarantees for peering (and perhaps other) traffic, and preserving sufficient headroom for bursting of research applications, can be comfortably met.
- Participation in a global collaborative around peering is strategic for the U.S. R&E community and is encouraged.

Introduction

The Future Peering Service Framework Subcommittee (Framework Committee) (list of members in Appendix A), under the leadership of co-chairs Pat Christian (University of Wisconsin) and Cas D'Angelo (Georgia Institute of Technology/SOX), presents the enclosed report for the development of the next evolution of Peering Services for Internet2 members. Since September, when the group was charged with creating a Future Peering Service Framework by the Network Architecture Operations and Policy Program Advisory Group (NAOPpag), this group had sixteen conference calls as well as a two-day face-to-face meeting and has discussed peering topics including:

- Scan of the Peering Environment
- Review of the Financials of TR-CPS
- Review of the Current Operations of TR-CPS
- Review of Current TR-CPS Architecture including Peering Location Details and Interconnect Utilization
- Planned and Conducted a Survey of Peering Practices by Regionals

During the face-to-face meeting, the Framework Committee agreed on the following principles as being vital to ensuring that the next evolution of a peering service continues to meet the needs of the Internet2 community. These principles include:

- Virtualization of the network rather than using a separate infrastructure
- Allowing NET + services as part of the peering service
- Creating a new Headroom Practice
- Continuing to offer dedicated TR-CPS ports
- Providing at least 24 months assurance of TR-CPS availability before planned major changes
- Increasing transparency and governance
- Addressing dedicated staffing
- Expanding collaborations to include international peering

Based on those principles, the Framework Committee provides recommendations on the key aspects for implementing the next generation of the Peering Service, with architecture, operations, governance, staffing and financial considerations.

It is the recommendation of the Framework Subcommittee, after consideration and deliberation, that the NAOPpag charge the NTAC in conjunction with Internet2 staff to engage the community to participate in developing and implementing the next evolution of a peering service for the Internet2 community and our global peers.

Strategic Value of the Service

Today, all major Internet Service Providers (ISPs) interconnect with one another and agree to exchange traffic at public and private peering locations. Most of these exchanges are "settlement-free", i.e., the parties pay only their own incremental costs to

interconnect one with the other. The Research and Education networking community has an established history of seeking settlement-free peering. Among the regional networks, peering has been seen primarily as a cost-saving measure. Whenever a regional network could exchange traffic with another large traffic source (whether another regional network or content providers such as Google or Microsoft, that may have a presence locally), the costs of doing so were much less than paying a Tier-1 ISP to carry the traffic for them.

National R&E networks also sought to peer with others, both regionally and nationally, for example, the Department of Energy's ESNet, Internet2's Commercial Peering Service (CPS) and/or CENIC and Pacific Northwest Gigapop's TransitRail (TR). Then, with CPS and TR combining their respective national R&E peering networks to form TR-CPS in 2008, further cost efficiencies were gained through economies of scale and consolidation.

However, settlement-free peering is not free. Peering has a cost in the personnel needed to maintain the network and peering relationships. The cost of the inter-connection equipment, colocation (space & power), riser fiber, peering exchange ports, etc., all contribute to the total cost of maintaining and operating a peering service.

Peering is about relationships
Building and maintaining personal relationships with other peering personnel among the community of content providers and national networks is surprisingly important. The relationships are built on a trust between networks that results in an improved level of service. Once that level of relationship and trust has been achieved, the responsiveness with which network issues are resolved is accordingly increased.

Traffic management and resiliency
One could look at managing multiple peering connections to the Internet much like that of managing a diverse financial portfolio. Bandwidth and demand should be balanced so that a disruption in any one service does cause undo harm to the overall connectivity portfolio. Peering provides opportunities to spread (and thus reduce) risk, particularly when peering through multiple locations. This is especially true for our R&E networking interests internationally. The ability to exchange traffic directly with our sister networks in Africa, the Pacific Rim and Europe at multiple peering locations only increases our mutual performance and independence from commercial pathways.

Overall, these factors all contribute to the value proposition of settlement-free peering to the R&E networking community. As our regional and national networks continue to evolve and grow, the development and support of this peering fabric should remain focused on the unique needs of R&E.

Goals and Objectives for the Next Evolution of the Internet2 Commodity Peering Service

At the highest level, the goals and objectives of the next evolution of peering service for the Internet2 community can be summarized in a straightforward manner:

1. A ubiquitous service available to all Internet2 constituencies
2. A cost effective service
3. A flexible service which can be used effectively by regional networks and higher education institution members (i.e., local policies can be accommodated)
4. A service that provides value for commodity use of the network without jeopardizing the research and educational (R&E) mission of the community (i.e., it will augment the R&E components and not compete with them)

One of the most important conclusions of the Framework Committee is that the value of a peering service is much more than direct cost savings. That is, while it is clear that the next evolution of the peering service needs to be cost effective, there is a great deal of value in defining the service as a highly reliable, highly flexible and robust service, or more correctly, a group of services, that meet the production needs of the Internet2 community. For example, in addition to commodity Internet traffic and content provider traffic, NET+ services seem appropriate for inclusion in the service.

Other Considerations and Goals—Near Term - Twenty Months

While architectural aspects of the peering service are described elsewhere in this document, in the near term, the next evolution of the peering service will target the following goals and objectives:

1. Virtualize the network – the network will leverage the existing 100 Gbps infrastructure and technically stable solutions to implement the service.
2. NET+ will be offered as a part of the service.
3. Headroom Practice – A more coherent approach to capacity and costs is to set data driven goals to assess those needs and costs so the community can make informed decisions. A "prime directive" of do no harm is mandatory for all types of traffic.
4. Flexibility – The service will offer dedicated ports for commodity services as well as the more traditional converged services currently offered.
5. Predictability – At least 18-24 months notice of any changes to the program's structure will be provided to allow regional networks and member schools to make adjustments to their operations and budgets.
6. Establish a default Autonomous System Number – The goal is to set a default "Higher Education" AS number to enable content and service providers "easy" access to the HE domain within the United States. (See next section for global ASN)

Other Considerations and Goals—Farther Term - Three to Five Years

The farther term considerations and goals entail initially establishing a coherent cost structure for the service. While the costs currently are incremental when compared to the ongoing total cost of operating the backbone, it is clear that a high value service must be sustainable in the longer term. Although the current 100 Gbps Internet2 network backbone provides the resources today for the peering service, the community cannot

realistically predict the future capacity needs of research and commodity in an age of big data. While some are able to articulate needs of their respective domain sciences, there is no coherent analysis which one may perform to obtain a reasonable statistical estimate. The community does have some aggregate data and snapshots in time of specific phenomena. Internet2 is in the process of ramping up its traffic analysis capabilities via deep data flow inspection. In conjunction and complementing the headroom practice, the future network requirements should be more data driven than casual estimates based on few criteria.

Dedicated Ports

The Framework Committee recognizes the need for dedicated ports. While the need for Advanced Layer 2 Service (AL2S) in research domain is clear, it is also clear that following the principle of using common production protocols and best practices, articulated elsewhere in this document, Internet2 should offer the ability for connectors to directly connect into the layer 3 production service for access to the peering service, without relying on AL2S for connectivity.

Service Guarantees and Headroom Practice

In the next evolution of the peering service, peering traffic should be protected from other classes of traffic from the connector handoff to the peer handoff so that the bandwidth committed to the customer (currently 20 Gbps of peering traffic per 100 Gbps interface) is always guaranteed. Thus, service guarantees should be extended to the connector backhaul links on AL2S.[1] The Internet2 headroom practice currently aims to preserve significant idle bandwidth on a given link to accommodate bursting by data-intensive research applications. The current version of the practice states that once a predictable traffic level on a specific link reaches 25-30% of the total link bandwidth, the link should be augmented with a parallel link on the same route segment. Given the added complexity of supporting multiple services over AL2S, the headroom practice should be revisited so that the contradictory requirements of providing service guarantees for peering (and perhaps other) traffic, and preserving sufficient headroom for bursting of research applications, can be comfortably met. The committee recommends that, as a starting point, 50% of the link bandwidth on the largest of a connector's backbone and connector link be reserved to accommodate bursting for research traffic.

Global Peering Service Opportunity

National Research and Education Network (NREN) partners in Europe have approached Internet2 about first interconnecting current continental peering networks and then later working towards a global peering service. While some aspects might be similar to current TR-CPS, the new program could integrate global investments and would have scale and efficiencies that would benefit not only Internet2 members, but also members of other NRENs elsewhere in the world. By acting together, the initial indications are we would further enhance domestic commodity peering, while also potentially continuing the development of infrastructure for global content initiatives. The aggregation of this key set of global peering traffic and the hardening of high-capacity interconnects between

[1] Backhaul service on the previous Layer 3 service was implemented by 10 Gbps lambdas, and thus guaranteed full use of all available bandwidth.

peering networks would also position the US R&E community for greater global NREN collaboration on other services, including intercontinental capacity alignment, and end-to-end research support.

Given global Internet traffic patterns that still rely on substantial hosting within the United States, Internet2 believes there is an opportunity to more widely share costs with global partners while also enhancing scale of peering services. There are also opportunities for global leadership coordination that could create new opportunities for US R&E community peering leaders.

The Framework Committee views participation in such a global collaborative as strategic for the U.S. R&E community and encourages the new oversight group to engage directly in oversight of these activities with the new peering coordinator.

Architecture Principles

The Framework Committee has developed a set of principles intended to guide the development of a comprehensive architecture and implementation of the service. These principles are intended to be the basis for a charge to the Internet2 Network Technical Advisory Committee (NTAC) and staff to aid in that development.

While the scope of these principles addresses only the peering service directly, it is recognized by the committee that in order to evolve a coherent overall architecture for the Internet2 network that includes peering as a service, it will be necessary to consider the interactions among all of the services that use the Internet2 network infrastructure.

Virtualization
One of the goals of the next evolution of the peering service is to efficiently integrate the peering service with the existing Internet2 network infrastructure rather than using the physically separate networks that are used for TR-CPS and the Internet2 R&E networks currently.

The architecture of this evolution of the peering service should include considerations for the best utilization of a single set of Internet2 backbone routers to implement both the R&E network and the peering service utilizing virtualized routing functions. The Framework Committee agreed that most of the concerns about changing to a virtualized infrastructure could be addressed by adhering to reliability and cost effectiveness principles. The architecture should also be flexible enough to support additional services as needed, such as Net+ or a high performance research service, etc. The committee discussed utilizing separate Virtual Routing and Forwarding (VRF) instances for each service, but defers the decision on actual design to the NTAC working in conjunction with staff. Additionally, the committee agreed that making the decision to virtualize on a single hardware base led to the conclusion that a single operational organization is needed to manage the network.

Transport
The architectural design should specify transport between peering routers that meets the reliability, scalability and flexibility requirements outlined in the following sections. While it is expected that the transport used in the short term will rely on separate Layer 1 circuits (as the current TR-CPS service does), the architecture should accommodate moving that traffic to the common AL2S platform in the future.

Reliability
In order for the peering service to function as required, reliability is a paramount consideration. The community, including the regional networks and members, expect the service to perform at production quality, with an availability approaching 100%. The expectation of reliability is also required for the service to be a member in good standing of the peering community. Outages or other poor performance will erode the trust needed to maintain the partnerships that make the service valuable to the community. The current architecture of the Internet2 network relies on AL2S as the underlying transport for higher-level production services such as Layer 3 backbone links and connector backhaul. AL2S also supports user-programmable services that are controlled via the Open Exchange Software Suite (OESS). While providing significant advantages in terms of flexibility and efficiency, AL2S carries some risk that its lack of maturity may negatively impact the peering service.

In defining the peering architecture, every effort should be made to ensure isolation between services in such a way that the reliability of the peering service is not impacted by other services. The peering traffic should be protected from competing with other classes of traffic from the connector handoff to the peer handoff so that the bandwidth committed to the customer (currently 20 Gbps of peering traffic per 100 Gbps interface) is appropriately prioritized. The interconnections between the Advanced Layer 3 Service (AL3S) and the peering service should be diverse enough to guarantee a high-level of resiliency in the case of link or node outages. There should be no single points of failure. Similarly, paths between the peering service and the connector should be diverse, and should allow connectors to multi-home to at least two locations on the peering service.

Scalability
The peering service should be designed with the capability to grow as the needs of the community expand. The expectation should be that the significant increase in bandwidth demand that the community has seen over the past decade will continue unabated. Currently, a connector is allowed 20 Gbps of TR-CPS bandwidth within each 100 Gbps AL3S service purchased. The new service should be designed such that Internet2 can increase that ratio as needed to support connector demand. Regardless of the ratio, connectors should have the capability to burst above their guaranteed committed bandwidth on TR-CPS, with packets above the committed rate delivered on a best-effort basis. Interconnections between AL3S and the peering service should be designed to allow for incremental growth (or shrinkage) at specific peering points as demand from the community and the availability of peering opportunities change.

Flexibility
The peering service should be flexible enough to meet the diverse needs of the community today and in the future. The committee discussed allowing regional networks to have the option of meeting Internet2 at one or more peering locations and connecting directly, or of backhauling the peering service over AL2S or a separate lambda as well as the possibility of accommodating a mixture of direct and backhauled connections on a per connector basis. The Framework Committee agreed that further community discussion on the topic was warranted and agreed to defer the final design to the NTAC working in conjunction with Internet2 staff.

The architecture should accommodate a connector's need to split their contracted peering service committed rate across multiple 100 Gbps ports as required. For example, a connector with two ports should be able to apply a 30 Gbps commitment to one port and the remaining 10 Gbps to a second port. Similarly, a regional networks should be able to apply all of its' commitment to one port if it so chooses.

In addition, the architecture should support the use of multiple 100 Gbps backbone lambdas per route segment. This should allow for flexibility in provisioning to accommodate the competing demands of service guarantees between the peering service, NET+, etc. It also would allow separation of traffic based purely on service type, if needed. For example, on links with heavy backbone usage it may be desirable to deploy a separate 100 Gbps to support only connector backhaul.

For the interconnection between the peering service and AL3S, the architecture should allow for changing local requirements for each peering location. For example, if a desirable peer became available at a different peering location within a metro area, the architecture should accommodate local expansion as required.

The provisioning flexibility described above should be designed in such a way that adjustments to configurations can be easily made as the demands of the community change over time.

Transparency
The peering service architecture should support flexible levels of reporting and visibility into the service's inner workings. Flexibility is required in order to provide information in the appropriate level of detail to the appropriate audience. Due to the sensitive nature of some aspects of peering, it is important to allow for the peering service operators to finely control the level of detail provided to each audience.

The architecture should support both low-level monitoring of important functions such as interface counters, as well as higher-level service monitoring that captures the availability and responsiveness of specific peers. Monitoring information should be made accessible via appropriate visualization and analytics capability in order to support both operational needs and reporting to management and peering service governance. Continuous and on-demand end-to-end testing should be supported. From an operational perspective, both monitoring and test tools should support operational responsiveness (particularly

important when dealing with peers and other peering service operators) and planning for continuous improvement of the service.

Optimization
The heart of the peering service resides at several geographically distributed peering locations. The current TR-CPS configuration in each of these metro areas has evolved over time based on the idiosyncrasies of local exchange points, historical opportunity, cost, etc. As part of the architecture and design of the new service, it is important to consider how best to optimize the peering service configuration in each of these locations, within the parameters that are under the community's control.

The architecture should support the ability to utilize common co-location and local interconnection opportunities where possible. In cases where local metro service is required to connect peering opportunities in different parts of the metro area, those services should be designed with the same levels of reliability, scalability, flexibility and transparency as described above.

Collaboration
The architecture of the peering service should support greater integration with the connector community's suite of services. The idea of a service-oriented networking approach has been discussed widely in the community of late,[2] in an attempt to provide a more integrated end-to-end network service that can support Net+ and similar initiatives.

In this spirit, it should be recognized that connectors might require some flexibility in the peering service definition in order to optimize the blended service they deliver to their customers. The service should accommodate configuration optimizations based on a particular connectors private peering and caching strategies, as an example.

Besides growing within the current geographic bounds, the peering service should be designed to support greater cooperation with international partners, particularly other NRENs that have already established peering services of their own. The committee sees this as an important objective that will support the needs of those in the community that are expanding or collaborating internationally. One example of this is the many higher education institutions that have established campuses outside of the continental U.S.

In addition, a peering service with international reach can attract a richer set of peering partners than are available to TR-CPS today.

The architecture should support a staged approach to integrating with international peering partners. The first stage envisioned would be the ability to share co-location space with a partner peering service, and exchange routes. The final stage would be a fully integrated peering service, with a single AS number[3]. The architecture should consider what intermediate stages might be required, and support all stages.

[2] See the reports from the NTAC face-to-face meeting in 2013.
[3] Note that this would likely require an organizational change to manage and utilize the single AS.

Operations

Peering Coordinator Roles and Responsibilities
The Lead Peering Coordinator (LPC) role requires that the individual be community-oriented and have the ability to deftly make decisions within the framework defined by the Peering Advisory Subcommittee (PAS). S/he would serve as the face of the peering service to the Internet2 community, as well as to commercial peers with whom s/he would engage. The LPC would need to have national peering experience, with global experience preferred, and must have a network of contacts in the peering community established. The committee suggests building a staff consisting of a full-time LPC and one or two junior peering coordinators as a transcontinental team. The LPC is required to identify peering and content distribution trends and set strategy for the peering service. It is critical that the LPC have experience in capacity planning. The LPC would need to maintain a healthy working relationship with Internet2 leadership and operational staff, the NTAC, the peering advisory subcommittee, and the steering group of fellow peering coordinators. The LPC will be responsible for community coordination regarding the service, and will serve as an evangelist for the service.

Advisory group/governance/oversight
The committee agreed that operations could be handled by a single entity that would operate under the guidelines provided by the PAS under the auspices of the NAOPpag. The PAS would define policy and peering strategy to provide a framework in which the LPC would work. The peering advisory subcommittee would employ a light-weight advisory process allowing LPC to provide a "heads up" on upcoming events or concerns to the PAS. The PAS would also review performance metrics and offer input as to where long term investments and improvements might be made. The PAS would be made up of 5-7 members of the community with individuals having a mix of expertise in peering and network architecture, CIO's and regional network executives and, where possible, would have geographic diversity.

The committee realizes the importance of transparency in the operation of the peering service, yet also notes that this must be balanced with the ability of the LPC to be nimble and able to take advantage of beneficial peering opportunities. It was noted that the LPC must have the ability to establish personal relationships within the peering community, and commitments s/he makes should not be undermined as that will cause him or her to be less effective. To this end, the committee suggested the creation of **a steering group** made up of two or three peering practitioners from the community who would serve as a "sounding board" for the LPC. This group would be available to provide community input to the LPC on tactical decisions for which the LPC seeks input. The LPC must have the ability to make the best decision based on his or her judgment, but the steering group would need to be made aware of tactical efforts that would benefit from community scrutiny. The LPC will provide a slate of candidates for approval by the PAS.

Dedicated Staffing
Dedicated staffing is required to ensure the best performing global peering network.

Global peering coordinators need to be responsive and have the knowledge to handle issues 24x7. The tight-knit peering community will quickly make changes to sideline peers that cannot respond quickly in an informed manner. Since so much of peering requires cultivating relationships, this requires dedicated resources to make sure that the best peerings are available to the community via the Internet2 network. Split resources will not have the time to give the attention and responsiveness required for the peering network to be successful. The committee suggests maintaining the existing budget for up to three FTE including the LPC and one or two junior coordinators reporting to the LPC who are focused on peering coordination activities - either employed by Internet2 or via alternate arrangements. These three FTEs would comprise a Peering Coordination Group (PCG).

Technical Support
Internet2 will provide technical support for operating the peering service, with a combination of support staff from the Internet2 NOC, and the PCG. This group will work closely with the greater NOC staff, which supports the larger set of Internet2 network functions.

Clear operational roles will be established by the Vice President of Network Services to ensure that the efforts of the Peering Coordination Group, the R&E IP network engineering team and the NOC are well coordinated and complementary one to the other. As NET+ grows, joint efforts will also need to be made by both the NOC and the peering coordination group to accommodate these needs. This process should be lightweight, as not to encumber the peering group or NOC in execution of necessary tasks. Underlying infrastructure will remain under the control of the Internet2 NOC and Network Planning Team.

With the move to an IP-VPN model for layer-3 services, the Internet2 NOC will also need to re-assess the roles of their staff in dealing with formerly innocuous configuration and maintenance tasks. The peering coordination group will assess requirements for operating in a VPN environment, and work with the Network Planning Team members in the NOC to establish a set rules for operating different elements of the base IP/MPLS network. This will include items such as IGP, iBGP, and MPLS configurations.

Maintenance events that will affect multiple services on the Internet2 infrastructure will be coordinated by the NOC, as they are now. The PCG will work with closely with the NOC to define the process for maintenance events and notifications.

Monitoring will be handled by the NOC, as it is today. The PCG will develop processes for handling outages, as well as configuration and maintenance tasks, that may vary from the procedures for the rest of the network. The nature of the peering network may require the additional tools to be used, or developed. The PCG will be responsible for developing the requirements, to be implemented by the NOC. With any new tools developed, the effort should be made to make them public, and the data accessible.

Ticketing will be handled by the NOC, as it is today. The assignees from the Service Desk will initially go to members of the PCG, who may delegate tasks to other members of the NOC staff as deemed appropriate.

A process will need to be developed for managing overlapping maintenance and operational requirements between the R&E and peering networks. In a shared environment, the services will need to be responsive to the needs of the user base, realizing that they are not entirely overlapping groups, and each serve different goals.

Capacity planning will be handled by the PCG. Periodic assessments will be relayed to the community, with data provided in a clear and transparent manner. The peering coordination group will work hand in hand with the Network Planning Team in the NOC to handle any necessary upgrades or augments to the underlying infrastructure.

International peering support
Among the international peering collaboration partners, such as NORDUnet, LPC and engineering contacts will be designated. These may be the same, or different, individuals, and will work on planning with the Internet2 PCG on relevant tasks. Periodic meetings will be held to address current and upcoming issues. The Internet2 PCG will handle any outcomes that require implementation on Internet2 infrastructure. Read-only access to the Internet2 peering routers should be granted to the designees of the international peering partner(s), to assist with troubleshooting and planning. Additional efforts will be made to share traffic and capacity planning information that is not already publicly available.

TR-CPS Financial Review

Internet2 and CENIC have compiled a shared view of the TR-CPS financials for the past four years. The financial detail includes the direct costs of the program, but like other Internet2 network-dependent programs, the generalized use of the Internet2 backbone, staff and overhead are not directly charged to this program. Included in the costs are 100% of the CENIC operating contract costs, the Indiana University service desk costs for TR-CPS and the costs of colocation, power, local loops when not provided by Internet2/CENIC and cross connects to the TR-CPS peering racks. The costs also include the maintenance, depreciation, and capital costs associated with the equipment used exclusively in TR-CPS.

Services	2010 - Internet2/CENIC	2011 - Internet2/CENIC	2012 - Internet2/CENIC	2013 - Internet2/CENIC	2014 - Internet2/CENIC
5050 - Assigned Employees	185,206	252,753	260,293	270,575	407,400
5225 - Travel	3,416	9,084	15,900	12,305	30,000
6210 - Legal	966	-	-		
6275 - Shipping	135	2,323	285	477	500
6602 - Backbone/Waves	-	-	18,000	523,069	523,069
6606 - Cross connect	588	142,082	200,628	196,561	179,828
6612 - Commercial Peering	450,560	697,617	582,858	556,374	814,126
6617 - Phone	2,000	1,867	1,632	1,615	1,452
6620 - NOC	131,250	257,000	277,082	292,420	347,947
6621 - NOC Travel	1,834	5,100	-		12,000
6650 - General	3,648	64,007	49,946	6,706	600
6661 - Maintenance, Software	-	16,650	6,850	8,800	2,400
6662 - Maintenance, Hardware	6,896	99,607	124,599	181,530	205,768
6663 - Maintenance O&M	-	-	186		
6672 - Colo, Racks	-	108,756	87,151	64,338	60,720
6674 - Colo, Power	-	73,759	78,385	84,811	75,251
N/A - "TR-CPS Support" Contra	(281,954)	(284,662)	(114,460)	(60,957)	-
6610 - Depreciation	46	166,496	244,930	251,815	255,561
Total	**504,590**	**1,612,438**	**1,834,264**	**2,390,439**	**2,916,621**
***Capital**	***224,039***	***1,117,642***	***57,001***	***18,823***	***840,000***
Total Expense <with Capital>	**728,629**	**2,730,080**	**1,891,265**	**2,409,262**	**3,756,621**

During 2014, it is expected that substantial capital expansion will occur for TR-CPS, with new higher-density and 100 Gbps capable core switches deployed to at least three TR-CPS peering cities. This will result in a substantial, but temporary, increase in overall costs, particularly due to depreciation of new hardware in the short term. At the same time, adjustments to the management approach and potential hybridization of the network hardware with other Internet2 services may result in some operational savings as a new approach may have some efficiency.

If the same methodology currently used for assigning direct costs to TR-CPS account and shared costs to Internet2's general network accounts is continued, the potential consolidation of the TR-CPS and R&E networks into a shared virtual network would mean that the direct costs to TR-CPS would be substantially reduced but the general costs to Internet2 would likely stay essentially flat (or reduce slightly). In a consolidated approach, costs associated with routers, core interconnects, power, colocation and other

support for the core routers may shift from being assigned to TR-CPS to Internet2's general network funds. (This consolidation may have an equally sized offset increase in the Internet2 general network budget.)

Consolidation of R&E and TR-CPS offers an opportunity to consider how many routers are necessary in each key R&E and TR-CPS geographic location. While consolidation might provide substantial operating savings over time, it is equally possible that scale or resiliency benefits of having multiple routers in key cities will continue to make sense. We would suggest the committees that will look at architecture consider expense as one of several important factors and that Internet2 be asked to report out potential consolidation cost implications as part of the analysis.

Metrics/Success Criteria

The Framework Committee reviewed metrics and success criteria with the Internet2 staff. A list of currently connected metrics along with suggestions for additional metrics to be collected is included in Appendix X. The Committee deferred final decisions regarding development of metrics to the NTAC.

Recommendations

The Future Peering Service Framework Committee recommends that the NAOPpag charge Internet2 staff to engage community participation via the Network Technical Advisory Committee—NTAC along with Internet2 staff in implementing the next generation of peering services for the Internet2 community and our global peers based on the architectural principles articulated by the Framework Committee:

20-month goals

1. Virtualize the network – the network will leverage the existing 100 Gbps infrastructure and technically stable solutions to implement the service
2. NET+ will be offered as a part of the service
3. Headroom Practice – A more coherent approach to capacity and costs is to set data driven goals to assess those needs and costs so the community can make informed decisions. A "prime directive" of do no harm is mandatory for all types of traffic
4. Flexibility – The service will offer dedicated ports for commodity services as well as the more traditional converged services currently offered
5. Predictability – At least 18-24 months notice of any changes to the program's structure will be provided to allow regional networks and member schools to make adjustments to their operations and budgets
6. Establish a default Autonomous System Number – The goal is to set a default "Higher Education" AS number to enable content and service providers "easy" access to the HE domain within the United States. (See next section for global ASN)
7. Staffing and Oversight --

a. Creation of a Peering Coordination Group consisting of Internet2 dedicated staff--Lead Peering Coordinator and two other staff members
b. Creation of a Peering Advisory Subcommittee (five to seven community members) would define policy and peering strategy to provide a framework in which the Peering Coordination Group would work.
c. **Creation of a steering group** made up of two or three peering practitioners from the community who would serve as a "sounding board" for the Lead Peering Coordinator.

Three-Five Year goals

- Future network requirements should be more data driven than casual estimates.
- Direct connections to the peering service should continue to be offered,
- Headroom practice should be revisited so that the contradictory requirements of providing service guarantees for peering (and perhaps other) traffic, and preserving sufficient headroom for bursting of research applications, can be comfortably met.
- Participation in a global collaborative around peering is strategic for the U.S. R&E community and is encouraged.

Appendix A
Future Peering Service Framework Subcommittee Members

Pat Christian, University of Wisconsin (co-chair)
Cas D'Angelo, Georgia Institute of Technology/SOX (co-chair)
Jeff Bartig, University of Wisconsin
David Crowe, University of Oregon/Oregon Gigapop (through October, 2013)
Dan Jordt, University of Washington
John Moore, MCNC
Jørgen Qvist, NORDUnet
Dave Reese, CENIC
Paul Schopis, OARnet
Michael Sinatra, ESnet

Appendix B
Metrics

Currently collected metrics:
BGP peer state
Interface status
Interface traffic statistics
BGP RIB updates
Aggregate traffic classification by type (eg, CDN, Video)
Total offered traffic load

Recommendations for additional data to be gathered:
Traffic stats (per-peer, and aggregate)

- Frame-size distribution
- IPv6
- Protocols

BGP update frequency, per peer
BGP route-refresh frequency, per peer
BGP prefix count, per peer (trend, set appropriate max-prefix automatically)
BGP prefix-count, aggregate for service
BGP peer stability (combination of peer state, update/refresh freq., prefix stability)

Appendix C

Interim Report of the Future Peering Service Framework Committee

On December 11-12, 2013, the Future Peering Service Framework Committee gathered for a face-to-face meeting in Ann Arbor, Michigan. Under the leadership of co-chairs Pat Christian (University of Wisconsin) and Cas D'Angelo (Georgia Institute of Technology/SOX), the committee spent two full days developing principles for the next evolution of Peering Service for Internet2 members. Since September, when the group was charged with creating a Future Peering Service Framework by the Network Architecture, Operations and Policy Program Advisory Group (NAOPpag), this group has met nine times and has discussed peering topics including:

- Scan of the Peering Environment
- Review of the Financials of TR-CPS
- Review of the Current Operations of TR-CPS
- Review of Current TR-CPS Architecture including Peering Location Details and Interconnect Utilization
- Planned and Conducted a Survey of Peering Practices by Regionals

The face-to-face meeting was the culmination of the fact-gathering efforts and the commencement of writing a report to be provided to the NAOPpag. While there is much work to be done in crafting the report, the committee wanted to share with the community a summary of the key areas where progress was made. This summary does not provide insight into the details of how the enhanced peering service will operate, but rather articulates the principles that the committee agreed would be vital to ensuring that the next evolution of a peering service continues to meet the needs of our community.

The key areas of progress include:

- Virtualization of the network—not separate infrastructure
- NET + should be offered as part of the peering service
- Create a new Headroom Practice—as a positive incentive to address the 20G "cap"
- Begin offering dedicated TR-CPS ports
- Provide at least 24 months assurance of TR-CPS availability before planned changes
- Increase transparency and governance
- Address dedicated staffing
- Expand collaborations to include international peering

Virtualization of the network - not separate infrastructure

The committee discussed the feasibility of changing from the current practice of using physically separate networks for TR-CPS and the Internet2 R&E networks to one that would utilize virtualized routing functions. For the current service, the community had preferred having a clearly separate infrastructure. The committee discussion converged

on an agreement that most of the community's concerns about changing to a virtualized infrastructure could be addressed by adhering to reliability and cost effectiveness principles. Further, the committee agreed that by using proven (not bleeding edge) technologies, such as those widely used in the commercial peering space (e.g., VRFs with MPLS), Internet2 could offer a reliable service on a shared, virtualized infrastructure. Additionally, the committee agreed that making the decision to virtualize on a single hardware base led to the conclusion that a single operational organization is needed to manage the network. The group indicated that these considerations may also allow the service to be delivered in the most cost effective manner.

NET + should be offered as part of the peering service

The committee agreed that the community wants clarity as to how NET+, TR-CPS and R&E routing will be handled and made progress in articulating a strategy that would allow maximum flexibility for all connectors. The group further agreed that, with the utilization of a common virtualized routing infrastructure, Internet2 could "rebrand" the services and state that there are multiple virtual networks (e.g., TR-CPS, R&E and potentially a very high-performance network for >10G end users) while also allowing each connector to choose how they wish to connect. The committee also agreed that it is critical to ensure that the proper "knobs" are in place and functional to provide all connectors with the ability to determine how they accept or advertise routes. (This is available today and the importance of ensuring that this capability remains was emphasized.) Next steps would include:

- Broader conversation with the community about the realignment and potential "rebranding" to move this forward
- In-depth conversation with the routing working group about implementing the principles suggested by the committee

Create a new Headroom practice -- as a positive incentive to address the 20G "cap"

The committee discussed community concerns about the current 20G cap on TR-CPS for each 100G link to Internet2. It was noted that the 20G cap was intended to provide an incentive for connectors to maintain sufficient headroom on their 100G connection for research flows (as well as help manage the capacity of the TR-CPS infrastructure). The committee thought it wise to plan in the future to deal with headroom issues directly. After discussion about the need to regularly keep the entire advanced networking ecosystem in sync around utilization, headroom, cost and equity of Internet2 services, the committee recommended that rather than set a cap on the amount of TR-CPS traffic that can be carried on a link, a new headroom practice that would ensure there was enough capacity to meet the "bursting" needs for purely research traffic be established both for the Internet2 backbone and for connections to connectors. The committee discussed using a starting point of ensuring there is at least 50G of capacity available for research traffic.[4] The group discussed options of requiring 50G of capacity be available on a

[4] In this case, research traffic refers to what has been considered by the community to be traffic that is in support of discipline research, and is usually generated from the campus. While we understand that some of this traffic may now come from commercial cloud sites (e.g., Amazon AWS), such traffic

single 100G circuit (not 50G spread across two circuits) and that each connector must establish routing such that research traffic can access that headroom. This model also considered that a connector with two 100G ports, might need to allocate 50G of headroom on only one of their connections. The discussion recognized that such a model could work as long as the community continues to monitor equity, cost and aggregation among connectors.

In the short term, until the next evolution of the Peering Service is implemented, the committee also discussed implementing a practice in 2014 in which Internet2 will monitor and measure the amount of capacity utilized by the regionals on each 100G link.[5] This would allow Internet2 to develop a policy that supports continued TR-CPS usage growth and ensures adequate capacity is reserved for research use. An annual process of allowing growth for the next year while also reviewing its impact and adjusting business models if needed might be considered.

Continue offering dedicated TR-CPS ports

The committee discussed continuing to support the ability of members to obtain a direct TR-CPS port on a peering router. The model discussed would use the existing flat rate fee for the port, which would not include the cost of transport from the connector to the port location. This would allow members to utilize existing infrastructure they may have to reach the peering exchange site, or choose to procure the appropriate transport service from Internet2 to gain access to the peering exchange site. The group discussed having such a port priced at the same fee as a port that includes Layer 3 access today ($200,000/yr) and monitoring over the next 18 months to determine if such a fee could sustainably fund the cost of TR-CPS growth among members. Further, the committee indicated that there are only a few regional partners who could place a large load on direct TR-CPS ports in the short term, therefore using the established 100G port fee is reasonable for the next 18 months.

---

would likely be considered NET+ or the "e" part of R&E Layer3 traffic, not what we have considered traditional research traffic.

[5] It is the intention of Internet2 to be as flexible as possible regarding the 20G cap for TR-CPS traffic. No regional will be penalized or charged additional fees if they exceed the 20G cap from time to time; as long as such bursts are not detrimental to the overall service. In fact, such bursts would be beneficial to Internet2 in determining growth in demand.

Provide at least 24 months assurance of TR-CPS availability before planned changes
The committee discussed the fact that many regionals need certainty in how the Peering Service will operate to make decisions on their other commodity Internet contracts or long term capital expenditures. The committee agreed to the principle that Internet2 will provide at least 24 months notice prior to implementing any changes that would degrade or significantly impact the business or operational models of the Peering Service.[6]

Increase Transparency and Governance
The committee agreed that operations can be handled by a single entity, separate from the governance but connected to governance. The committee further indicated that it is critical that governance be open and transparent as well as having community oversight. The committee recommended engaging the community in defining policy and governance and that a Peering Advisory Group be formed at the request and under the auspices of the NAOPpag. Agility and credibility are important aspects of any governance process and, therefore a "light-weight" advisory process was suggested allowing the Peering Coordinator (see next section) to provide a "heads up" on upcoming events or concerns to the Peering Advisory Group. The Peering Advisory Group would also review performance metrics and provide input to Internet2 in setting strategic direction.

Address Dedicated Staffing
The committee outlined the responsibilities and qualities necessary for the role of Peering Coordinator to support this effort. It was agreed that a dedicated Lead Peering Coordinator would be needed. The Lead Peering Coordinator would be responsible for capacity planning, coordination with Internet2, and community coordination & evangelization and, due to the importance of this role, the Lead Peering Coordinator needs to be dedicated to Internet2 for peering. The group also indicated that it will be important for the Peering Coordinator to be able to maintain credibility within the larger national/international peering community.

The committee further recommended that a small team be built to include one or two additional junior peering coordinators. Team members could come from the community and would provide input on matters at a tactical level. This team will be subordinate to the Peering Advisory Group. It was also felt that a very small group of peering coordinators throughout the community should be created (with rotating seats) to provide guidance to the Peering Coordinator.

Expand Collaborations to include International peering
The committee agreed that extending the peering infrastructure internationally would bring more peers to the enhanced peering service and would allow additional collaboration with international partners. Internet2, CENIC and Nordunet have already

[6] Of course, there may be times when extenuating circumstances might prevent Internet2 from adhering to this principle – but in such cases the community will be informed and engaged in discussion on the best way for the Peering Service to proceed.

begun conversations about a North American/European collaboration and this conversation will be expanded and pursued. (Jorgen Qvist of Nordunet is an active member of the committee and was present in the Ann Arbor meetings.) The belief is that the global R&E community may find benefit through continued collaboration and unified presentation of a future service offering.