



**Implementing the Layer3 Future Peering Group  
Vision:  
Internet2's Interpretation and Response**

**V1.1  
21 April 2015**



# Table of Contents

<b>Table of Contents</b> .....	<b>3</b>
<b>Background</b> .....	<b>4</b>
Purpose .....	4
Current Architecture .....	4
<b>Future Architecture</b> .....	<b>5</b>
Platform Consolidation .....	5
TR-CPS Backbone Headroom Principles .....	6
<i>TR-CPS Capacity Management Principles</i> .....	7
<b>Net+ Traffic Integration</b> .....	<b>8</b>
Net+ Background .....	9
Net+ Future .....	9
<b>Network Management</b> .....	<b>9</b>
Transparency.....	9
Peering Steering Group .....	9
Operational Data .....	10

## Background

### Purpose

This document is intended to outline the steps and policies that the Internet2 engineering team have developed in response to (one or more of) the Future Peering Committee's recommendations for the TR-CPS service. In brief, the major goals, as outlined in the report, were:

- 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

The focus of this plan is on the immediate changes needed to support the Future Peering Committee's recommendations on the next 6-12 month horizon. The document has been reviewed and augmented by members of the Layer3 Community Group, which includes several members of the previous Future Peering Committee, and the Layer3 Community Group endorses the plan of action developed by Internet2 staff and outlined in this document.

### Current Architecture

The existing Internet2 TR-CPS service exists in a compartmentalized fashion, on separate routing hardware and separate backbone infrastructure. The seven TR-CPS Juniper MX960 routers are largely within the same cities as the 10 R&E routers, but are often in separate colocation facilities. These seven facilities were chosen because of their proximity to the nation's largest commercial peering infrastructures. They are interconnected by one or more 10 Gbps Ethernet waves on Internet2's optical footprint. Connectors generally access their TR-CPS services on the same edge port that they access R&E services. The TR-CPS traffic is carried over a point to point network circuit to multiple TR-CPS routers for redundancy.

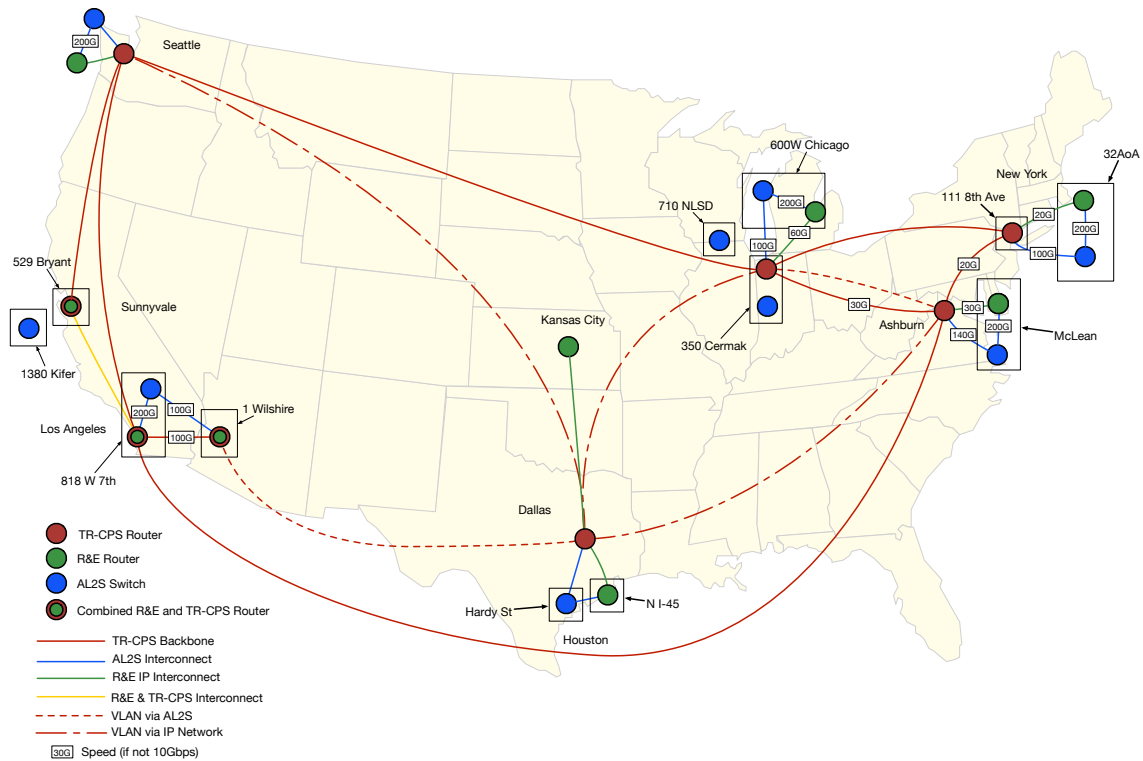


Figure 1: Current TR-CPS Infrastructure

Connectors receive their R&E and TR-CPS routes on separate peering sessions with separate autonomous system numbers. For comparison, the R&E network typically carries approximately 14,000 routes, while the TR-CPS network is closer to 250,000 routes. In addition, Internet2 carries a small, but significant, number of commercial routes in its R&E routing table in support of its Net+ program. These routes, limited to entities that have entered into a Net+ agreement with Internet2’s members, are distributed to the Connectors via the R&E peerings, but not the TR-CPS peerings. In some cases, the Net+ entity may present routes to Internet2 in both the R&E and TR-CPS tables- in which case the Connector can expect to get both sets of routes.

## Future Architecture

### Platform Consolidation

It’s clear that both the TR-CPS infrastructure and the Research and Education infrastructure serve different needs. While the R&E network is

focused on abundant bandwidth, large elephant flows, and the maintenance of vast amounts of latent headroom, the TR-CPS network typically serves smaller flows and is more focused on maintaining an infrastructure investment profile that more closely mirrors the predictable traffic profile. The economics of the two networks are also vastly different with the R&E network serving traffic exchange between two Connectors with postalized network fees, while the TR-CPS network facilitates exchange between Connectors and commercial networks that requires periodic adjustment and investment. It's also the case that the Connectors largely treat these networks differently from both a cultural and operational stance. For these reasons, Internet2 will continue to maintain a logical separation of the two networks for the foreseeable future.

However, there are gains to be made by physically integrating the two networks. Internet2 will merge the TR-CPS and R&E logical routed instances onto the same hardware platform. In some cases, this will eliminate redundant equipment, and in other cases, it will grow the reach of the two networks beyond their current footprints.

To accomplish this, Internet2 will configure the existing R&E routers to maintain an instance of the TR-CPS network on the same router, and vice-versa. The R&E focused traffic will continue to traverse the AL2S network between the core nodes, while the TR-CPS focused traffic will continue to use the existing 10G backbone links that have been re-homed into the combined architecture. However, Internet2 staff may utilize the AL2S network for traffic growth in an opportunistic fashion. These decisions will be logged and reviewed by the Peering Steering Group.

With the consolidation of the two services onto a single platform, there will undoubtedly be times when the activities of one service might conflict with the activities of the other service. As the decision to prioritize one service over another is a particularly difficult one to craft in a unilateral fashion, Internet2 staff will use their best judgment with the individual cases where this comes into question. The decisions will be logged and reviewed in collaboration with the Peering Steering Group. Internet2 will also periodically revisit the overall question of service prioritization on a regular basis with the Internet2 Community. The timeframe for this revisit will be decided with the Peering Steering Group, but will not exceed an annual review period.

### TR-CPS Backbone Headroom Principles

The current (as of 2014) Internet2 headroom practice doesn't directly address the process for augmenting the TR-CPS backbone links. It's crucial that Internet2 staff maintain an appropriate level of headroom on the network to

accommodate gradual traffic growth, usage pattern changes due to temporary network outages, and large transient events (e.g. Apple software distribution, Netflix series availability, etc.). As a general rule, commodity network traffic is much more predictable than research and education network traffic and can be much better characterized with network analytics. In addition, the ecosystem of commodity network providers introduces a level of “art” to the management of network traffic in relation to the various peers. Finally, R&E networks are focused specifically on ensuring there is plentiful capacity to support bursty behavior from the research applications. Commodity networks are more commonly focused on the most efficient way to exchange traffic within the given ecosystem of commercial costs. This list of principles seeks to balance the need for deterministic and documented expectations of Internet2 staff with the need for discretion and contextual evaluation in the capacity management process.

#### TR-CPS Capacity Management Principles

- Internet2 should seek to minimize the need to transport TR-CPS network traffic between Layer3 core nodes. Where possible, traffic between Internet2 Connectors and TR-CPS peers should drain locally within the same core-node. Where Connectors and Peers are separated by the backbone, Internet2 should endeavor to transport Connectors to the TR-CPS node via the AL2S network. Internet2 will create a set of best common practices (BCPs) for TR-CPS Network peering that will help Connectors steer their traffic flows to the most appropriate TR-CPS router. These BCPs will contain- but not be limited to- configurations that seek to minimize TR-CPS backbone traffic.
- Internet2 should strive to maintain enough headroom for large-scale transient network events. Anticipated events will be shared with the community in advance, where possible, along with any augments or mitigation strategies Internet2 has put in place to address the event. These events will be logged and reviewed monthly by the peering steering group.
- Internet2 will maintain per-peer network statistics that support the analysis of usage patterns between network peers and Internet2 Connectors. This will generally be collected via flow data and be stored in accordance with Internet2 network management privacy policies. The flow data analysis methods will be aimed at identifying flow transactions at the organizational level. This data will be reviewed in a weekly report by staff. Peers with Private Network Interconnect (PNI) ports that are over 60% utilized will be flagged in the report and discussed on a weekly engineering call. Internet2 will also take into account the overall bandwidth exchanged with each peer when making augmentation decisions. Each peer discussed will be logged with the decision for monthly review by the peering steering group. The peering steering group reviews will be logged with the decision.

- Internet2 will maintain a soft bandwidth limit of 20Gbps for each 100Gbps layer 2/3 port under contract. Internet2 will not maintain any hard rate limiting on the edge connections and that bandwidth may be spread across the Connector's edge ports. However, Internet2's recommendation that each Connector maintain at least 50% of their bandwidth on each network interconnect for research and education traffic remains intact. Internet2 will maintain per-Connector network statistics. This will be reviewed in a weekly report by staff. Connectors that are utilizing over 60% of their allotted TR-CPS bandwidth will be flagged in the report and discussed on a weekly staff engineering call. Each participant discussed will be logged with the decision for monthly review by the peering steering group. The peering steering group reviews will be logged with the decision. As a courtesy, Internet2 will make the Connector aware that its usage is over 60% of its allotted TR-CPS usage.
- Internet2 will endeavor to utilize the Internet2 Layer1 network to augment bandwidth between Layer3 core nodes for TR-CPS backbone traffic. Internet2 may, at its discretion, utilize the Internet2 AL2S network for *inter-node* traffic growth for opportunistic needs. Internet2 will maintain statistics on aggregate backbone usage. Engineers will review a report on backbone usage weekly and the peering steering group will review the report monthly. Internet2 will log all decisions to augment backbone capacity or utilize AL2S and all new decisions will be reviewed monthly by the peering steering group.
- Peers on a common exchange point will be monitored and reviewed weekly by Internet2 staff. These peers will be flagged for review when they consume a more than 1.5Gbps of the common exchange interconnect transport. Decisions regarding these peers will be logged and reviewed monthly by the peering steering group.
- In an effort to pro-actively manage the expected demands and latent capacity of each TR-CPS subscriber, Internet2 staff will coordinate closely with the Connectors to understand their anticipated needs and traffic models. This information will be captured and regularly reviewed for traffic engineering optimization and peer-facing capacity modeling.

## Net+ Traffic Integration

One of the goals of the Layer3 integration is to provide greater flexibility on how Internet2 Connectors receive their Net+ routes. In addition, the combined infrastructure provides additional opportunities for connectivity that should be discussed. This section runs through some of those considerations.



## Net+ Background

Internet2 currently connects with all Net+ peers via its R&E network. Those routes are only advertised to its Connectors via the R&E peerings. Internet2 has contractual relationships with Net+ providers to ensure that their routes are distributed to the Internet2 Connectors. Internet2 is aware that some Connectors don't peer with the TR-CPS network or- more commonly- don't send the full set of their customer routes onto the TR-CPS network. Given the focus of the R&E network support of large network flows, the R&E infrastructure is generally configured to support large Net+ providers such as Amazon Web Services via its jumbo-frame-enabled 100G backbone. This extends into the Connector networks as well, with many Connectors treating their R&E connections more openly than their TR-CPS connectivity.

However, there are scenarios where connecting a Net+ provider to the TR-CPS network may be desirable. It could be easier to peer with smaller Net+ providers via the commercial IXPs that are homed to the TR-CPS AS. Internet2 also understands that some Connectors find it desirable to obtain connectivity to commercial Net+ peers over their infrastructure that's dedicated to commercial connectivity.

## Net+ Future

To address these needs, Internet2 will entertain peer-facing connectivity in both the R&E and TR-CPS logical networks. It will also make the routes available to Connectors across both sets of peerings. There are some implementation details that need to be worked out regarding the incongruent set of routes that both networks receive from the Connectors. These will be identified and discussed publicly within the Internet2 NTAC. Internet2 will provide Connectors with the controls needed to manage their preferred method of receiving Net+ routes.

## Network Management

### Transparency

Internet2 will continue to maintain an operationally transparent stance. All regular reports, notes and decisions will be openly logged and made available to the wider Internet2 community. For confidential information, Internet2 will log as much of the information that can be made public.

### Peering Steering Group

Internet2 will convene a small, focused advisory group that receives regular and frequent updates on peering requests, architectural changes, operational issues, and augmentation needs. They will receive more detailed reporting on

the state of the network core and edge and will participate in the process of applying the commodity headroom principles. While the peering steering group is not meant to be a body that will review confidential information, it may be necessary that some items will need to remain internal. However, in keeping with the operational transparency principles outlined above, all reports that don't contain any confidential information will be available to the wider NTAC.

### Operational Data

The consolidation of the network will also drive the consolidation of many of the Internet2 NOC's management tools. There will be some behinds-the-scenes enhancements that will allow for a more holistic and unified approach to capacity planning with the combined infrastructure. Internet2 will work with the Peering Steering Group to evaluate its reporting tools to ensure maximum usefulness to both Internet2 staff and its Connectors.